Bolt Beranek and Newman Inc.

(12)

Report No. 4527

AD A121429

Movements Information Network (MINET) Testbed Design Study

D. Hunt, P. Sevcik, B. Hitson, and K. Pogran

December 1980

Prepared for:
United States European Command
and
Defense Advanced Research Projects Agency

82 11 16 037

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER	2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) Movements Information Network (MINET) Testbed Design Study	5. TYPE OF REPORT & PERIOD COVERED Design Study
	6. PERFORMING ORG. REPORT NUMBER 4527
7. AUTHOR(s) D. Hunt, P. Sevcik, B. Hitson, K. Pogran	8. CONTRACT OR GRANT NUMBER(s) MDA903-78-C-0129
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 50 Moulton Street, Cambridge MA 02238	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order 3491
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Blvd., Arlington, VA 22290	12. REPORT DATE December 1980
	13. NUMBER OF PAGES 156
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Defense Supply Service - Washington Room 10 245, The Pentagon Washington, D.C. 20310	15. SECURITY CLASS. (of this report) UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

UNCLASSIFIED/UNLIMITED

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

computer networks, Movements Information Network, MINET, ARPANET

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

The MINET concept is an application of state-of-the-art networking technology to the needs of the cargo movements community of the U.S. armed services in the Western European theater. The objective of this study has been to obtain MINET user requirements, propose an architecture for the MINET testbed, prepare a plan for operation, maintenance, and future growth of the testbed, and provide budgetary cost estimates and a risk analysis.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Report No. 4527 Bolt Beranek and Newman Inc.

MOVEMENTS INFORMATION NETWORK (MINET)
TESTBED DESIGN STUDY

D. Hunt, P. Sevcik, B. Hitson, K. Pogran

December 1980

Submitted to:

Headquarters
United States European Command (USEUCOM)

and

Director
Defense Advanced Research Projects Agency (DARPA)

A

EXECUTIVE SUMMARY

This summary highlights the conclusions reached in the final report of the Movements Information Network (MINET) Testbed Design Study. The objective of the study has been to obtain MINET user requirements, propose an architecture for the MINET testbed, prepare a plan for operation, maintenance, and future growth of the testbed, and provide budgetary cost estimates and a risk analysis. The MINET concept is an application of state-of-the-art networking technology to the needs of the cargo movements community of the U.S. armed services in the Western European theater.

The MINET concept was developed by representatives of the U.S. European Command Headquarters (HQ USEUCOM) and the armed services. The preliminary functional requirements for the MINET were published in the Statement of Requirements (SOR) document, dated 2 March 1980. This report describes a design that can meet the requirements of the SOR. It represents a minimum risk application of proven networking technology in a peacetime testbed environment. The testbed will be used to evaluate the utility advanced networking techniques have for the movements community.

The major recommendations and conclusions reached during the MINET Testbed Design Study are listed here. Each of these recommendations and conclusions is discussed in depth in this report.

1. Use ARPANET technology for the MINET testbed.

   This recommendation is based upon the SOR, together with considerations of cost-effective operation and readily available technology. Relying on ARPANET technology allows the testbed MINET to be deployed with the lowest possible risk.

2. The MINET testbed should be a separate network.

   The MINET testbed will be built as a separate network in the theater and connected to the ARPANET with internet gateways. Initially there will be a single gateway to the ARPANET.

3. The MINET testbed will have 114 terminals at 80 sites.

   There will be 80 military sites that will be connected together with the testbed MINET. These sites will all be in the European theater, and will be able to access movements data bases both in Europe and in the United States. A total

of 114 terminals will be distributed among these 80 sites.

4. The MINET testbed will use 16 voice-grade trunk circuits.

Sixteen trunk circuits are adequate to support the projected traffic flow in the testbed network. With the exception of the Europe-CONUS link, each of the IMPs is at least doubly connected to the network, so the net can tolerate the failure of any single trunk circuit except the CONUS circuit. Analog voice-grade telephone circuits will be used for the network trunks, with modems that can provide a digital data rate of 9.6 Kbps per second. The majority of these circuits will be leased from European common carriers (PTTs), and, where available, others will be provided by DCA Europe.

5. The MINET testbed can be built using 13 IMPs and 12 TACs.

A testbed network with 13 IMPs (Interface Message Processors) and 12 TACs (Terminal Access Controllers) is sufficient to support the estimated traffic load of the 114 terminals. The IMPs are packet-switching minicomputers that connect to MINET trunk circuits as well as to MINET host computers; a TAC is a specialized host, connected to an IMP, which serves as a terminal concentrator. Each IMP, except for the one in the Europe-CONUS gateway link, will have a collocated TAC.

6. The testbed will give fault-tolerant, responsive service.

The performance goal of the testbed MINET is to provide an overall network delay of less than 2 seconds. The reliability goal is to meet the performance goal for at least 90% of the user terminals, under any probable single component failure. The analyses in this report indicate that the testbed will meet these goals, and that the performance goal can still be met even if there is a fourfold increase in the projected amount of testbed network traffic.

7. The MINET testbed will have three electronic mail hosts.

Three EM (electronic mail) hosts are recommended for the MINET testbed. They will be in Frankfurt, Naples, and London. These hosts will support the UNIX operating system, as well as an electronic mail system designed to meet the needs of the testbed MINET users.

8. The Network Control Center will be in Europe.

The Network Control Center (NCC) for the MINET testbed will be located at Camp King, Frankfurt, the headquarters of the 4th Transportation Brigade. The ARPANET NCC, at BBN in Cambridge, Massachusetts, can serve as a backup NCC for the MINET testbed. Each NCC is a host computer that is dedicated to displaying network status, statistics gathering, fault isolation, and remote testing.

9. DCA Europe is responsible for acquisition and operation.

DCA, the Defense Communications Agency, has been designated as the organization that will be responsible for the acquisition and operation of the MINET testbed. DCA will not only assume responsibility for the acquisition and maintenance of the MINET trunk circuits, but also for the acquisition, siting and operation of the testbed network hardware.

10. The testbed will be deployed in three stages, over 4 years.

The MINET testbed will be deployed in three stages. In each stage a portion of the total network will be implemented: Stage 1 will include most of the sites in West Germany and the UK; Stage 2 will include most of the sites in the Mediterranean area; and Stage 3 will include sites in Greece and Turkey, plus a few others. Stage 1 sites are planned to be on the network one year after initiation of the testbed MINET development effort. Stage 2 sites are planned to be on the net two years afterwards, and Stage 3 sites are planned to be on the net three years after development begins. Thus only in the fourth and final year of the testbed program are all sites planned to be on the network.

11. The MINET testbed estimated budget is 14.5 million dollars.

The budgetary estimate for the MINET testbed program is 14.5 million dollars, in constant 1980 dollars. The final project budget estimates, for the purpose of DoD planning and funding requests, are being developed by the U. S. European Command (USEUCOM). USEUCOM has further developed the budget estimate given in this report, by including certain contingency costs and the costs for a System Engineering and Technical Advisory (SETA) contractor. These cost revisions yield an overall estimated program cost of 16.0 million (constant) dollars.

1. INTRODUCTION

This report is the final report of the Movements Information Network (MINET) Testbed Design Study, which has taken place between March and December of 1980. The objective of the study has been to obtain MINET user requirements, propose an architecture for the MINET testbed, prepare a plan for operation, maintenance, and future growth of the testbed, and provide budgetary cost estimates and a risk analysis.

1.1 Project Scope

The concept of a Movements Information Network (MINET) is an application of state-of-the-art networking technology to the needs of the cargo movements community of the U.S. armed services in the Western European theater. The MINET concept was developed by representatives of the U.S. European Command Headquarters (HQ USEUCOM) and the armed services. The preliminary functional requirements for the MINET were published in the Statement of Requirements (SOR) document, dated 2 March 1980. This report describes a design that can meet the requirements of the SOR. Although the major SOR requirements are summarized in this report, this report does not supersede the SOR.

1.1.1 Report Plan

This report is divided into five major sections. The first section provides an introduction and a project perspective. The second section takes the requirements of the MINET users and maps them into requirements for the network itself. The third section takes the requirements for the network and maps them into a specific network design that satisfies those requirements. The fourth section suggests how to obtain the necessary circuits and hardware for the network design, how to operate the network, and presents a project plan for the testbed development, installation, and operation. The fifth section assesses risks associated with the MINET development, and related research and development projects.

Without loss of continuity, the reader can skip parts of this report and gain an understanding of what the MINET is, and what costs are involved. Section 1, the project perspective, is recommended to all readers. Section 3.1 describes the recommended MINET testbed configuration, sections 4.1, 4.2, and 4.6 describe the plans for the network circuits, hardware, and manpower respectively, and section 4.7 contains the budgetary cost estimates for the testbed.

1.1.2 Purpose of MINET

The purpose of the MINET is to improve the managing and tracking of cargo movements into and within the European theater. This improved capability applies to the movement of cargo from ports of embarkation (POE) in the continental U.S. (CONUS) to ports of debarkation (POD), and from PODs and other in-theater origins to controlling movements regions and to selected activities.

The first step towards realizing MINET will be to develop a testbed network, which will provide HQ USEUCOM as well as the headquarters of the three armed services (USAREUR, NAVEUR, and USAFE) and their component commands with an opportunity to evaluate the chosen MINET networking technology in the European theater and to assess the effectiveness of this technology as a communication medium between logistics data bases in CONUS and data bases in the theater. The objective of the testbed approach is to implement a representative portion of the longer-term network. This portion includes enough geographically dispersed users so that it can be of immediate use, and so that it can provide a non-trivial basis for evaluation. On the other hand, the testbed is limited in scope and cost in order to hasten the completion of the evaluation and in order to keep the evaluation focused on essential MINET features.

Ultimately, the MINET should be able to serve the movements community in the theater in a wartime situation. The testbed MINET, however, will not be capable of providing the degree of robustness that will be required in wartime. Over the course of the next few years, the kind of technology that can provide this robustness will become more economical, and can then be incorporated into the MINET. The design for the testbed MINET presented here allows for evolution to a wartime-ready network while requiring only minimal changes to the MINET terminal user interface.

The testbed MINET network will also provide substantial improvements to the current peacetime operation. In particular, it will provide the movements community with real-time, rather than historical, status of cargo shipments. Responses to queries against the logistics data bases will be provided within minutes, rather than within a few days.

1.2 Overall Project Guidelines

In order that the testbed MINET be considered a success, it must be demonstrated that networking technology has become sufficiently stable that it can support the demands of the movements community in the European theater. Accordingly, the network should be designed to be a minimum risk network -- it should concentrate on the use of off-the-shelf technology. The

minimum risk guideline also implies that the technology should not be pushed too close to the limits of its capacity. For this reason, the number of user terminals attached to the testbed network should be carefully controlled. The use of off-the-shelf technology not only contributes to the minimum risk goal, but also helps reduce costs and hastens the deployment of the testbed network.

Implicit in the notion of a testbed is the potential for a complete, or at least enlarged, facility at some point in the future. Accordingly, the network should permit smooth growth of both capacity and functionality, and, over a wide range, the cost of adding an additional user should be constant. A basic consideration, therefore, is how much growth the chosen technology will support. It is important to understand what aspects of the technology might become bottlenecks in future growth plans, such as the desire to triple the number of users or to double network traffic. This growth should be able to be accommodated without any major modifications to the user interface to the network.

1.3 Approach

A designated set of representatives from the MINET user community, called the MINET Working Group, has been tasked with participating in the design, development, operation, and

assessment of the testbed. The Working Group is headed by representatives of HQ USEUCOM, and includes cargo movements specialists from each of the three services. The Working Group reports to the General Officer Steering Committee (GOSC), which also has representatives from all three services and HQ USEUCOM.

In order to obtain a careful analysis of the issues involved in building the MINET testbed, HQ USEUCOM requested that the Defense Advanced Projects Research Agency (DARPA) initiate a study, with DARPA serving as the funding agency and contract monitor for the study contract. The study has taken place between March and December 1980. The objective of the study has been to propose an architecture for the MINET testbed, prepare a plan for operation, maintenance, and future growth, provide budgetary cost estimates, and present an analysis of risk and related research and development efforts. The preliminary results of the study were presented to the MINET Working Group and the GOSC in August 1980. The GOSC has approved the MINET testbed concept and the budget estimates presented in the August briefing. HQ USEUCOM and the three services are presently planning the funding of the MINET testbed implementation, operation, and maintenance contract.

The first phase of the study was to survey a representative set of potential MINET users. This survey yielded two kinds of data: (1) a functional requirement, which expresses the needs of

the potential users, and (2) a traffic requirement, which shows how heavily the potential user community uses existing communication facilities to track cargo shipments. The next phase of the study reduced the traffic requirements data into bandwidth requirements between the various network sites. Next, the bandwidth requirements, taken together with well-established methods for laying out trunk circuits in a packet-switched net, yielded a family of possible testbed network topologies. Functional requirements were taken into account while laying out the packet-switched subnet as well as while considering the hosts to be attached to the subnet. Having proposed a testbed network topology, the study effort provides an O&M plan for the testbed, followed by budgetary cost estimates. Finally, the study has included a risk analysis and a discussion of future requirements and evolving technologies that can be tailored to meet them.

2. REQUIREMENTS

The first major subtask of the MINET design study is to map the requirements of the testbed network's future users into requirements for the packet-switching subnetwork and for the electronic mail hosts. To begin, then, it is necessary to identify the future users. The future users or their representatives should be able to provide both the users' functional and the users' traffic requirements. These can be mapped into network functional requirements and network traffic requirements respectively. Practically all of the user functional requirements can be stated without any knowledge of packet-switching technology, and with only a high-level understanding of electronic mail systems. The user traffic requirements can be stated independently of either the subnetwork or electronic mail characteristics. Accordingly, it has been possible to interview individuals with direct responsibility for day-to-day communication but who may have little exposure to the current networking technology. In this section, the user requirements are summarized and network requirements are derived. The user traffic requirements are explicitly mapped into network traffic requirements.

The initial list of testbed MINET users and the initial set of functional requirements appear in the Statement of Requirements document (SOR). We obtained additional functional

requirements during two site survey trips in March and June of 1980. Potential testbed MINET user sites were given traffic requirements forms prior to the June site survey. Both during the June survey and for several weeks thereafter, we received the completed traffic questionnaires that provided the basic traffic requirements for the MINET testbed.

We obtained final adjustments to the user requirements, as well as to the list of network users, during the MINET Workshop Meeting at Patch Barracks in August 1980. Participants in the MINET Workshop included the MINET working group and the General Officers Steering Committee.

2.1 User Functional Requirements

Functional requirements for the testbed MINET user community are deriv d from three sources: the SOR, the site visits to potential MINET user locations, and a design phase which included commentary and feedback from USEUCOM representatives. These requirements are summarized in the following subsections.

2.1.1 Requirements from the SOR

In this subsection, major user functional requirements from the SOR are summarized. In those cases where the SOR requirements have been amended by the Working Group, the amended requirements are given.

Since foreign nationals will need to access the testbed MINET on a routine basis, only unclassified traffic will be carried by the network. Nonetheless, the testbed design should include consideration of encryption methods appropriate for supporting unclassified traffic. A minimum requirement for the testbed is that it protect against hostile exploitation of aggregate logistics information transmitted through the network. In support of this requirement, major links within the network should be protected using cryptographic devices.

The testbed network is to be developed in four phases in order to control costs, assess network performance, and further evaluate MINET user requirements. The functions provided in each phase are summarized here, and described in more detail in separate paragraphs. Phase 1 provides a network of terminals that can communicate using a standardized electronic mail service. Phase 2 extends the phase 1 capabilities by allowing the terminals to query certain movements data bases in Europe and CONUS. In phase 3, the phase 1 and 2 capabilities are extended to support update of the data bases from the remote terminals. Phase 4 includes the capabilities of phases 1 through 3, and also supports direct transfer of data from one movements data base host to another.

The phase 1 testbed network will provide electronic mail service. This service will support real-time terminal-to-

terminal conferencing, and management inquiries associated with movements management and control. A major role of the electronic mail service is to demonstrate how movements control can be improved by access to more timely data.

In phase 2, testbed network terminal users will be able to submit queries to and obtain responses from host systems with movements data bases. Phase 2 will support on-line editing and verification of query forms in order to reduce errors and facilitate rapid response.

In phase 3, testbed network users will be able to update the movements data bases, subject to appropriate access controls. Phase 3 will include the capability for on-line editing and verification of data base update forms.

Phase 4 will support bulk transfer of information between movements data base hosts. Once a terminal operator has initiated the data transfer, the operation should continue automatically until the transfer is complete.

The SOR contains a list of sites that will have MINET terminals, as well as a list of the movements data base hosts that are to be accessible to the testbed MINET users. During the course of the design study, however, the set of MINET testbed sites and data base hosts changed several times. On August 15, 1980, the MINET working group agreed upon a testbed network that

would include 114 terminals at 80 sites, and which, starting in phase 2, would provide access to 4 data base hosts. The most recent list of sites and data base hosts are included in this report; the hosts are listed below and the sites appear in Appendix A.

It should be emphasized that a "site" is the finest-grained organizational subdivision that has appeared in the SOR or in subsequent discussions. Thus, for example, the Karl Schurz Kaserne in Bremerhaven has five MINET sites: (1) Headquarters, First Movement Region, 4th Transportation Brigade; (2) Traffic Management Office (TMO), First Movement Region, 4th Transportation Brigade; (3) Headquarters Military Sealift Command (MSC) Europe; (4) Military Traffic Management Command (MTMC), Transportation Terminal Unit (TTU), Bremerhaven; and (5) the USAFE Water Port Liaison Office (WPLO). Each site may have one or more terminals, and each terminal may have one or more users. The terms "site," "terminal," and "user" are useful in expressing the requirements of the SOR. Furthermore, application of communications technology typically groups sites into "nodes." For example, a telephone central exchange node serves a community of nearby user sites. The 80 testbed MINET sites have been grouped into 12 nodes, and the grouping of sites into nodes is shown in Appendix A. The justification for grouping the sites into these 12 nodes is given in subsection 3.1. In the above

example, the Karl Schurz Kaserne, which has five MINET sites, is a single MINET node. Using this terminology, the MINET testbed configuration has 12 nodes, 80 sites, 114 terminals, and some larger number of users.

It should also be emphasized that there are two kinds of host systems planned for the testbed MINET. The movements data base hosts, which will be referred to as "data base hosts," are to be made accessible to network users by phase 2. The electronic mail hosts, or "EM hosts," which support the electronic mail service, will be available in phase 1.

Of the four data base hosts which are to be accessible to MINET testbed users, one is in Europe and three are in the CONUS. The European data base host is DAMMS (Department of the Army Movement Management System) at Camp King in Oberursel, Germany. The U.S. data base hosts are the CAPS (Consolidated Aerial Port System) of the Military Airlift Command (MAC) at Scott Air Force Base, Illinois; the U.S. Army Materiel Development and Readiness Command (DARCOM) Logistics Intelligence File (LIF) at the Presidio, San Francisco; and the MTMC Eastern Area Terminal Management System (TERMS) at Bayonne, New Jersey. All of these hosts operate in an unclassified mode.

The testbed network facilities are planned to be available for use 24 hours a day, seven days a week, with the exception of

scheduled preventative maintenance. Operator support, however, will be available on a more limited basis. The SOR does not specify the operating schedules for the testbed data base host systems; clearly the CONUS hosts must be accessible in the early morning hours, local time, if they are to support inquiries from the theater during standard working hours.

All equipment used in the MINET testbed must be able to operate using 220 volt, 50 Hz. power sources. This will eliminate the need for external power converters. Although it is not required that the testbed hardware be Military Specification equipment, it should be resistant to normal shocks and vibrations that may occur while it is being moved or serviced.

2.1.2 Requirements from the Site Surveys

This subsection presents the major user functional requirements obtained during the site surveys. The requirements for electronic mail service, although obtained during the site surveys, are presented in the separate subsection on electronic mail.

Terminals on the testbed MINET will be one of four types: hard copy, display, portable, and bulk input terminals. Most sites preferred the hard copy terminals in order to keep local records of their transactions. Bulk input terminals, such as terminals with card readers, are required by some sites in order

to input manifest data directly to the data base hosts. Some sites have existing word processing terminals which, if possible, should be interfaced to the testbed MINET.

The site survey showed that the testbed users are interested in a formatted input capability. This will allow local editing and verification of queries and updates. Most users interviewed preferred to fill out a form on a terminal screen, which requires a display terminal. However, a prompting scheme, which will work on hard copy terminals, was also acceptable to those interviewed.

Access control facilities should be provided, both for the data base hosts and for the mail in the electronic mail system. Access control for the data base hosts must be provided by the hosts themselves. An authentication and access control method for electronic mail should restrict access to mailboxes to certain user groups or individuals.

A means of allowing testbed MINET users to communicate directly with Telex users will be most useful, since many locations involved in movements activities will not be MINET testbed sites, but will have Telex terminals. Testbed MINET users will wish to use the forms capability to compose and edit Telex messages prior to sending. This direct Telex-MINET connection can also allow Telex users to submit queries and updates to the data base hosts.

2.1.3 Requirements from the Design Study

In order to install the MINET testbed within a reasonable time frame and with lowest possible risk, the testbed must be based on ARPANET technology. The justification for this recommendation can be broken down into five sequential steps. First, a switched (rather than point-to-point) network is needed given the large number of hosts and users which must communicate with each other. Second, of the switched networks, packet switching, rather than message or circuit switching, is most appropriate. Message switching cannot provide the fast response needed for the MINET. Circuit switching would result in higher line costs and offers no statistical multiplexing advantage for carrying bursty traffic, while packet switching does. Circuits would be required to connect each site or cluster of sites to an electronic mail host, and this topology, unlike the packet-switched topology, is vulnerable to single line failures. Third, of the packet-switched networks, a dedicated net serving the military movements community is preferable to a leased packet-switched service provided by common carriers. The common carrier X.25-based services will not be available in all the MINET countries in the required time frame. Further, the PTT regulations restrict the options for gateway connections to other nets, and trans-border data flow. Finally, there is no off-the-shelf equipment for data protection over these networks.

Fourth, of the dedicated packet-switched networks, those developed by the U.S. Government are better suited to the MINET user community than those provided by commercial vendors. This choice allows the MINET to conform to the Department of Defense Standard Internet and Transmission Control Protocols, which are not supported on the commercial nets. Conforming to these protocols allows the MINET to take advantage of current developments in network security as they become available. Fifth, of the two possible U.S. Government-developed packet-switched networks, AUTODIN II and ARPANET, the preferred choice is the ARPANET. The ARPANET is a seasoned technology, while AUTODIN II is still under development. Furthermore, ARPANET technology emphasizes widely dispersed nodes, whereas AUTODIN II emphasizes fewer, higher volume nodes. Current AUTODIN II plans call for only two nodes in western Europe, which would force long access lines and result in lower overall reliability. The ARPANET model, consisting of numerous, lower volume nodes with nearby terminals, more closely matches the needs of the MINET users. For these reasons, the MINET Testbed Study has focused on providing a network that uses the ARPANET technology.

A further requirement, which did not appear in the SOR but which has been accepted by the Working Group and the General Officer Steering Committee, is that the testbed deployment be

divided into three stages. In each stage a portion of the total network will be implemented: Stage 1 will include most of the sites in West Germany and the UK; Stage 2 will include most of the sites in the Mediterranean area; and Stage 3 will include sites in Greece and Turkey, plus a few others. In terms of trunk circuit quality and availability of hardware support, the Stage 1 area is lowest risk, while the Stage 3 area is highest risk. Dividing the testbed deployment into these stages makes it possible to focus on the lower risk areas earlier in the project. Stage 1 sites are planned to be on the network one year after initiation of the testbed MINET development effort. Stage 2 sites are planned to be on the net two years afterwards, and Stage 3 sites are planned to be on the net three years after development begins. Thus only in the fourth and final year of the testbed program are all sites planned to be on the network.

Based upon the information received from the testbed MINET Working Group and the General Officer Steering Committee, the testbed user community will consist of 80 sites with one or more terminals at each site, for a total of 114 terminals. The total terminal count does not include some word processing terminals at some of the sites which may be used for access to the testbed MINET. There are 43 sites in Stage 1, 24 sites in Stage 2, and 13 sites in Stage 3. Each site, together with its terminal requirement, is listed in Appendix A.

Another requirement identified during the design study is to provide a statistics gathering capability within the electronic mail hosts. This will keep track of the quantity, length, number of addressees, and other information about electronic mail messages sent by the testbed MINET users. This information can be used to tune the operation of the network, and to serve as a basis for predicting usage patterns in a full-scale MINET.

2.2 User Traffic Requirements

Traffic on the testbed MINET comes from a number of diverse sources. Almost any type of communication mechanism currently being used can be supplemented, or in some cases replaced entirely by Electronic Mail (EM). In the following sections we identify the sources of traffic and how they can be converted into a common unit of measure for the purpose of comparison, estimate the total traffic volume and flow, and discuss the limitations of the data.

2.2.1 Sources of Traffic

Conversions from current communications methods will be the initial source of traffic for the testbed MINET. Included are AUTODIN, telephone (commercial, military, AUTOVON), Telex, and other manual methods such as courier and physical mail. Most of the Telex traffic and a large fraction of the AUTODIN and voice traffic could be moved to MINET EM.

2.2.2 Survey of Traffic Volume and Characteristics

Data for estimating the volume of traffic and its distribution in the testbed MINET was initially collected by a survey of the sites involved. In this survey, each site was asked to report the number of AUTODIN messages, Telex messages, and phone calls it would exchange with each of the other potential sites on the testbed MINET. We arranged the reported data in three categories: AUTODIN, Telex, and Voice (commercial phone, military phone, AUTOVON). A condensed version of the raw data that was used in establishing the traffic volume and patterns is included as Appendix B.

In addition to traffic volume and distribution, information was also obtained about the characteristics of the traffic. It was determined that an average telephone conversation lasted about 5 minutes. The percent of each type of traffic that could be converted to EM was determined to be 50%, 75% and 50% for AUTODIN, Telex, and Voice respectively. It was also determined that there are approximately 30 lines of text in a typical AUTODIN or Telex message. Half of the total traffic was expected to occur during the busiest two hours of the day (the same two hours for most users).

Some survey questions were not answered by many people and others produced a number of conflicting responses. Our best

estimates for these quantities are introduced in the following sections as they become necessary.

2.2.3 Conversion of Raw Data to Messages/Month

We begin by converting raw data from the survey into units of messages per month (msg/mo). By a message, we mean the amount of information that is transmitted by one user to another user at one time (e.g., one Telex or one AUTODIN message). Our survey of typical AUTODIN, Telex and EM messages on currently existing systems showed that they are all in the range of 1200-1500 characters/message. This agrees well with MINET user estimates of 30 lines per typical AUTODIN or Telex message. With 30 lines of 40-60 characters/line, we get 1200-1800 characters of data per message. When we include the additional characters required for the addressee line, subject line, message ID, date, time, and so on, we have typical message sizes of from 1400-2000 characters. To be on the safe side, we will use 2000 characters per message as a typical message size in all further calculations.

The data received by us was reported in a number of different units covering differing lengths of time. All AUTODIN and Telex traffic was converted to msg/mo by using 30 lines/msg and 20 working days/month as appropriate. All rounding is done to give conservative estimates (i.e., higher traffic).

The conversion from voice traffic to msg/mo is somewhat less exact. Given that we have a 5 minute phone conversation as typical, one must establish how many EM messages it will take to replace it. We can consider three main types of phone calls: calls to disseminate information, calls to get information from someone who knows the answer we want, and calls to discuss something with someone (conversational). The first type of call can be replaced by one EM message since no response by the callee is needed. The second type of call may typically be modelled as two EM messages, one to ask the question(s) and one to get the answer(s). The conversational type of call is much harder to convert. Perhaps one person raises questions which the other answers partially but then proposes alternatives or other questions that the first person then answers, and so on, all in the same phone conversation. This leads to at least two cycles of the "ask-answer" sequence. On the other hand, sending EM to another person might cause the sender to think more carefully about the questions or to group many questions into one call that might have previously required more than one call to answer. This would tend to cause fewer EM messages than we might expect from a direct conversion of telephone minutes.

The previous paragraph is only a partial discussion of the assumptions that must be made in converting from voice minutes to number of equivalent EM messages. Our purposes only require that

we have typical values for all types of phone conversations averaged together. For this, we will use 2 EM messages per 5 minute (or less) call. This should be close enough and, if anything, is probably on the conservative side.

2.2.4 Total Traffic Estimate

In addition to the potential sites discussed in the SOR, there were a number of terminals and sites that were added at the MINET workshop in Germany during the week of August 11-15. Due to the late date at which these were added, it was not possible to get detailed traffic data for these sites. By consulting with potential users at the August workshop, we determined that 200 msg/mo/terminal would be reasonable. This is the number that we used to estimate traffic for the 40 additional terminals.

Although our survey asked each site to report both number of messages sent and received, most sites did not report messages received. The few sites that did list messages received showed them to be identical to messages sent in almost every case. This is not surprising for Voice or Telex traffic: each of these is addressed to only one user or office at a time. AUTODIN, however, has the capability of multiple recipients which will cause more than one copy of a message to be received for each one sent. Furthermore, addressing to multiple recipients is typical for many EM applications as well. Given the lack of survey data

for messages received, as well as the differences in addressing modes between Voice and Telex on the one hand and EM on the other, we will use only the messages sent data from the user survey. The next portion of our analysis concentrates on messages sent (composed). Following this, estimated received message traffic, based on experience with existing EM services, will be incorporated into our traffic analysis.

The total estimated composed message traffic for the MINET is just the sum of the three types of traffic as described above. Applying the conversion factors discussed previously -- 50% of AUTODIN at 1 EM msg/AUTODIN MSG, 75% of Telex at 1 EM msg/Telex, 50% of Voice at 2 EM msg/5 minutes of phone call -- to the data in Appendix B gives 5709 msg/mo for AUTODIN traffic, 3110 msg/mo for Telex traffic, and 10724 msg/mo for Voice traffic. There is an additional 8000 msg/mo from estimated traffic on the 40 terminals that are not included in the SOR. Summing these numbers and rounding up to the nearest 1000 msg/mo gives a total estimated testbed MINET traffic of 27,000 msg/mo composed. This is the number we will use for standard peacetime operating conditions in the fully deployed (Stage 3) testbed MINET. Note that of the total, 30% is contributed by AUTODIN, 13% by Telex, and 57% by Voice.

A discussion of traffic during peak loading conditions, and of traffic during exercise conditions, can be found in section

3.1. Node-to-node flows showing exactly where the traffic originates and where it is sent are shown in Figure 2-1. The unspecified traffic flows of the 40 terminals added during the August meetings are entered in the last column marked with an asterisk. The nodes shown in Figure 2-1 represent the site grouping shown in Appendix A.

\RECD SEND\	LON	RDM	BRM	FRK	RAM	HDL	STT	NAP
LON	135				35			138
RDM			5	9	45	6	3	
BRM		111	393	601	570	105	82	
FRK	51	205	480	2252	774	538	161	
RAM	30	35	84	322	954	468	38	90
HDL		16	138	314	305	108	50	
STT		16	96	207	213	217		
NAP	148				90			
SIG	115							1445
ROT	118				25			895
IST					40			
ATH								
USA	95				88			165

Figure 2-1 - Node-to-Node

SIG	ROT	IST	ATH	USA	*	TOT	
28	35			43	1600	2014	LON
				6	1000	1074	RDM
				18	1600	3480	BRM
				275	400	5136	FRK
	25	40		437	600	3123	RAM
					400	1331	HDL
					200	949	STT
1195	1195			80	800	3508	NAP
	690			60	-200	2110	SIG
520				360	400	2318	ROT
					800	840	IST
					400	400	ATH
145	160					653	USA
					TOTAL =	26936	

traffic (Msg/mo) composed

2.2.5 Correlation with Previous Experience

To ensure that we have not made some severe procedural or arithmetic error, and that we have included all major sources of traffic, we would like to have some way to judge whether this number of 27,000 composed msg/mo is indeed reasonable. The data shown in Figure 2-2 is the result of a study of typical patterns of electronic mail use. It shows the number of messages sent and received (sent/recd) per day by a typical user on the vertical scale, and the amount of time he has been using EM on the horizontal scale. For the business environment (where "users" are "offices" which receive a lot of mail addressed to one or many different people), the numbers shown are typically somewhat higher than for individual users. The study indicates a 1 to 3 ratio of sent to received messages, for both kinds of users. A separate curve for business users and is also shown in Figure 2-2. For individual users, the maximum of 40 msg/day (10 sent and 30 received) is mostly limited by the time and resources an individual user can devote to sending/reading EM. The maximum of 80 msg/day for the business user is limited by the physical impossibility of reading and responding to more than this many messages on a single terminal in a given working day.

The range of the vertical scale goes from inexperienced users on the low end to experienced EM users on the high end. These two curves should serve as upper and lower bounds on the actual behavior of EM on the testbed MINET.

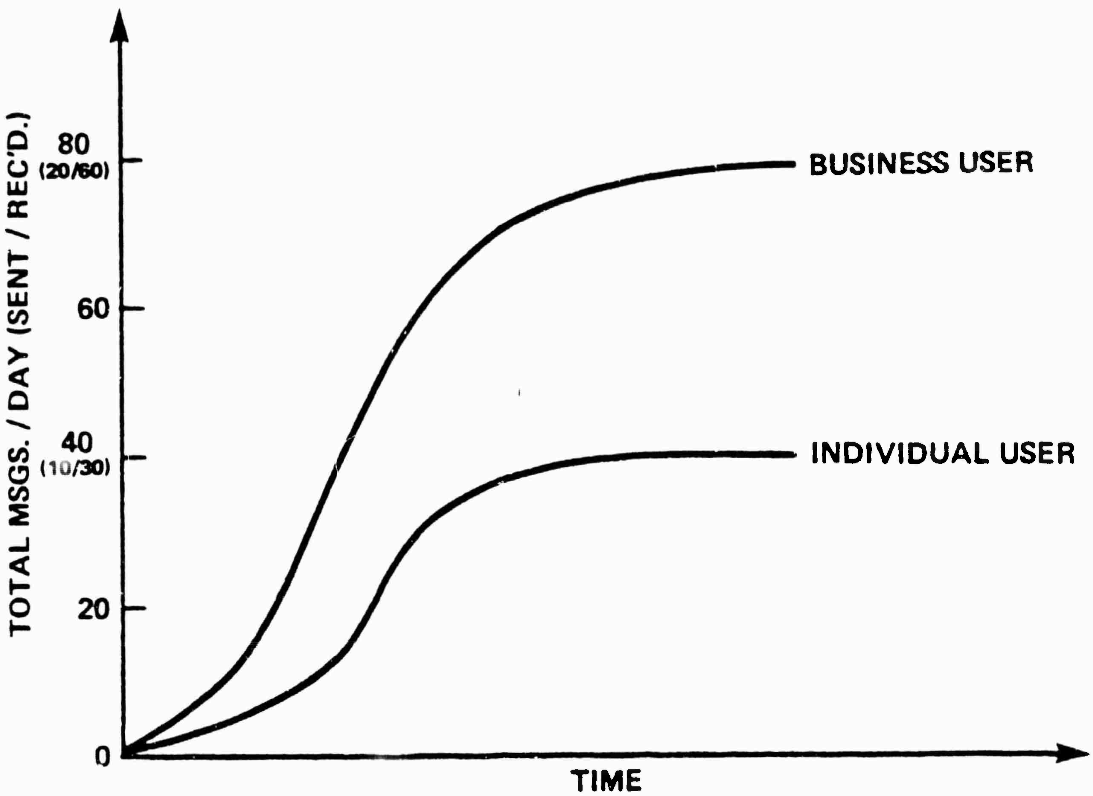

Figure 2-2 - Message Volume as a Function of User Experience

The total estimated traffic of 27,000 composed msg/mo, obtained from the survey, divided by 114 terminals gives us a value of 240 composed msg/terminal-mo. If there are 20 working days per month, this gives us 12 msg/terminal-day that are composed, or 48 msg/terminal-day that are both composed and received, using the 1 to 3 ratio of the study. We see from Figure 2-2 that this is a reasonable value for business applications, where there are many users per terminal. This

might indicate that the MINET environment is more closely related to a business environment than a single-user environment.

The ratio of 1 message sent to 3 messages received is based on the experience of EM users in the ARPANET and elsewhere, and is the best (if not the only) data available. Typical use might have a person send a message to one primary recipient with a copy to a second person and to himself. Thus, there is one message sent (composed) and three received (read), one each by the primary recipient, secondary recipient, and original sender. We will use the ratio of 3 messages received (read) for every 1 message sent (composed) to complete our estimate of total MINET traffic. The estimate of 27,000 composed msg/mo is not inconsistent with past experience and we will use it as a target number in further considerations for testbed MINET design.

Applying these conversions of 3 to 1 ratio for received traffic to the information in Figure 2-1 gives us the message sent and message received rates for each of the testbed MINET nodes, as shown in Figure 2-3. The data in Figure 2-3 indicate the messages per month sent by all terminals at each node, including messages sent to other terminals at the same node. Similarly, the figure indicates the messages received by all the terminals at each node.

Node Number	Node Location	Total Messages/Mo Sent	Received
1	London	2014	6042
2	Rotterdam	1074	3222
3	Bremerhaven	3475	10425
4	Frankfurt	5136	15408
5	Ramstein	3123	9369
6	Heidelberg	1331	3993
7	Stuttgart	949	2847
8	Naples	3508	10524
9	Sigonella	2110	6330
10	Rota	2318	6954
11	Istanbul	840	2520
12	Athens	400	1200
13	USA	653	1959
	Total	26931	80793

Figure 2-3 - Messages/Mo. Sent/Received at each Node

3. MINET SYSTEM DESIGN

This phase is a synthesis of the data contained in the requirements section and expresses communications requirements in terms of network functions and performance. This transformation from user related requirements to network design reflects not only user needs but also practical project considerations and constraints. An essential part of this phase is the verification of the requirements, scope and objectives, both for completeness and relevance to the user community and also to see whether the characteristics defined result in a practical network. The model includes physical assumptions, constraints and numerical parameters describing the network. This model is realized in functional elements required to support the network's features.

The design alternatives and architectural issues are presented at a relatively high level of abstraction since the detailed packet switching technology is predefined. The MINET will consist of a separate ARPA-like network which will be interneted to other networks. This design consists of tailoring the ARPANET technology for the specific MINET requirements. The foundations for the MINET system design are the functional and traffic requirements described in Section 2. The network will require data protection. Electronic mail will be provided as an integral service on the network. Movements information data base hosts will need to interface to the net.

3.1 Network Design

The overall goal of the network design is to provide a testbed on which the users can evaluate movements information transfer alternatives. The goal is not to test packet switching technology, but to test packet switching's utility to a specific segment of the military operations in the European theater peacetime environment. The network must, therefore, provide reliable and responsive data communication service within the constraints of the environment.

The first decision is the selection of a reliable data communication technology. The key driving factors are the final projected traffic per month and the geographic distribution. This generally focuses the technology options. A very modest traffic and geographic requirement would obviate the need for a special network, while a very large traffic and geographic requirement may require non-off-the-shelf solutions. Fortunately, the projected traffic can comfortably be handled by the ARPANET technology.

A network must fundamentally rely upon the quality and availability of transmission facilities. There are many types of data communication circuits available in parts of Europe, both from the PTTs, and from the U.S. Defense Communications System (DCS). The range quickly narrows upon the application of a few practical considerations:

- Digital circuits are available only in Germany. The digital synchronization technique of the German PTT, the Deutsche Bundespost (DBP), makes it impossible to use off-the-shelf interfacing hardware. These character isochronous circuits are particularly difficult to encrypt with reliable recovery from errors.

- Digital circuits are being installed on the DCS under the Digital European Backbone (DEB) upgrade program. The DEB is significantly behind schedule. Furthermore, many of the MINET cities are not in the DEB plan.

- Analog circuits are available to all of the cities identified either from the PTTs or DCS. However, the general quality of the circuits is poor outside of Germany, England, and the Benelux countries. Broadband circuits are generally not available.

- Voice grade circuits are probably the only circuits available during an exercise. It would be best to design the network to rely on the lowest common denominator circuits in order to have the greatest flexibility in exercises.

The conclusion is to lease voice-grade circuits with CCITT M1020 conditioning. These are the most generally available circuits. They are fully compatible with equivalent circuits

available on the DCS. The MINET can use very sophisticated modems in order to get as high a data rate as achievable. It is expected that a rate of 9600 bits per second (bps) will be consistently achievable throughout the network. The service quality parameters of M1020 provide further assurance that the data rate will be achievable. It must be noted, however, that some circuits may not sustain M1020 conditioning. During those periods, if the circuit is still operative, the modems should automatically downgrade their rates to 7200 bps, 4800 bps, and finally 2400 bps. This may need to be done by the packet switch controlling the modem across the DTE/DCE interface.

It would be best to rely on a mix of circuits from the DCS and PTTs. The disjoint failure modes of the independent systems will provide a more reliable overall operation. This is particularly true in local access to long haul transmission facilities. As an example, two PTT circuits leaving Patch Barracks would probably follow the same distribution system to the branch point for the long haul circuits. If one of the circuits were a DCS circuit, there would be no common elements. DCS circuits are, however, unlikely to be available for testbed use. We have therefore, planned and budgeted for using 90% PTT leased circuits.

The goals of reliability and responsiveness must, therefore, be provided in an environment of relatively (to current ARPANET

conditions) unreliable and slow circuits. This is the overriding system constraint. The SOR does not define specific reliability and responsiveness requirements which are appropriate for a packet switched network. The following design will provide the best economically feasible and minimum risk solution to meet the goals. Specific steps will be taken at each opportunity to optimize responsiveness and reliability. The final design will then be carefully evaluated and shown to provide acceptable projected service.

3.1.1 Node, NCC and EM Host Placement

Generally expensive and unreliable long haul circuits are managed extremely efficiently in packet switched networks. Packet switch nodes permit many subscribers to use the circuits seemingly simultaneously. Each subscriber uses the circuit only when data is being transmitted to or from the service hosts. The most cost effective placement of nodes is a difficult optimization problem which must account for the costs of multiplexed and direct access lines (dedicated per terminal), costs of trunk circuits (shared among many terminals), and the cost of nodes (node costs must be amortized over many terminals and must include maintenance). The range of solutions extends from a single large node (a star network) with many costly access lines, to a packet switch at each terminal concentration, which can eliminate all access lines.

This network has many other constraints which preclude a strictly economically-based solution. The terminals and users have therefore been grouped geographically so that a modest number of nodes can provide service with a minimum of relatively short access lines. The quantity of nodes is sufficiently great to preclude any large concentration of subscribers that would rely on a single node. The recommended 12-node configuration, shown in Figure 3-1, represents a practical point in the spectrum between a star network and a fully distributed network.

The quantity of nodes is large enough to subdivide the user population to the point that a typical node failure would affect less than 10% of the users. The quantity is, however, small enough to be a manageable facility for the purposes of a testbed.

Major clusters of users which are geographically co-located (within approximately 100 KM) are defined to be a node. The node includes a TAC and an IMP which can service up to 64 terminals, 4 hosts and 4 trunks. There are advantages to placing many nodes in a dispersed fashion close to the users:

- This minimizes cost of access lines. If each node has two trunks leaving it, then it becomes cost effective to cluster the cost of approximately three or more terminal circuits through a node (plus the cost of node hardware).

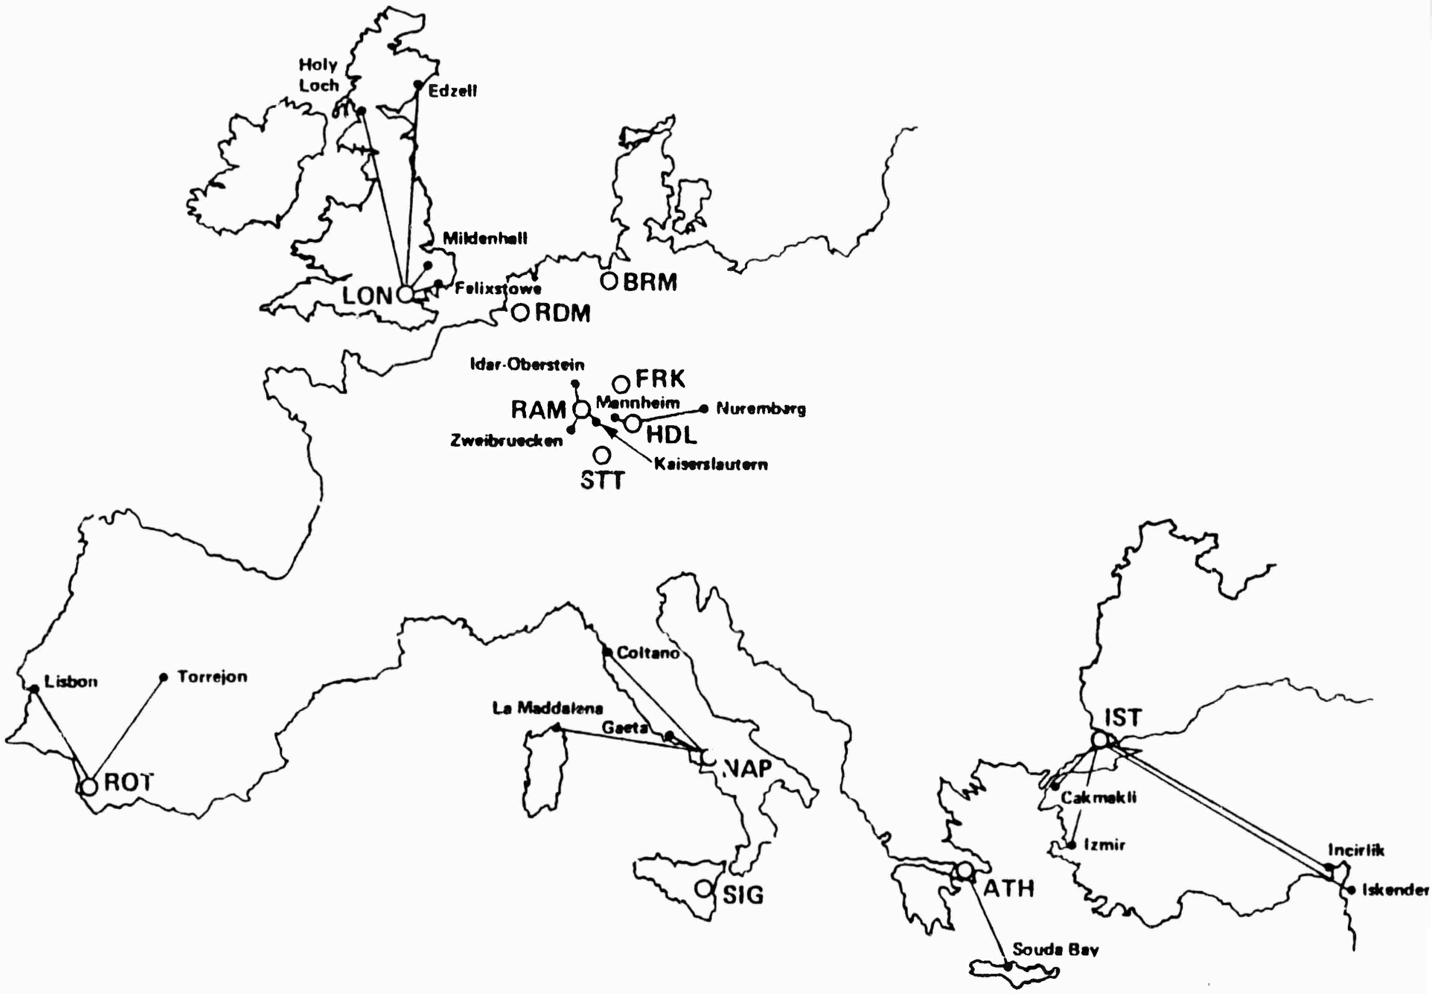

Figure 3-1 Testbed MINET Nodes and Remote Sites

- Packet switching provides reliable service through dynamic routing. The less non-packet switched service (dedicated access lines), the more reliability is presented to the users.

The NCC location is based upon practical network management considerations. The network should be managed from a single location. This location requires software and hardware maintenance staff along with training and operations personnel. The site should, therefore, provide physical support for the required personnel and hardware. It should also be centrally located with good transportation connections to the remaining sites. Finally, the site must be relatively easy to reach by airline from the U.S. in order to provide good spare parts delivery. The candidate sites are Stuttgart, London and Frankfurt. The site which was selected by the working group is Camp King, outside of Frankfurt.

The EM hosts are located using an iterative analysis which reduces the "cross network" traffic as much as possible (locations which generate greatest traffic) while providing a uniform load on the hosts. This process locates EM hosts at London, Frankfurt, and Naples. All terminals at those sites are considered local, while all other nodes must be homed to one of these sites.

The number and geographic extent of MINET sites prohibits single deployment of the network. The network deployment is divided into three stages which correspond fairly closely to the geographic regions covered. The division is based upon site criticality (Government provided) traffic requirements, and site technical risk (site maintainability and expected circuit quality) ratings. A weighted rating biased towards minimum risk provides the basis for the network staging.

3.1.2 Topology Design

The network must be sized to handle peak busy hour demands. The projected messages per month must be converted into packets and then to bits per second to and from each node during the busy hour. The conversion must account for both originated and read messages. This process will also generate the projected number of messages, packets and bits passing in to and out of the EM hosts in the busy hour. We can then use bits/sec (busy hour) to determine line capacities and node throughputs.

The total traffic reported earlier does not occur in a uniformly distributed way throughout a 24-hour day. Previous experience and user survey responses indicate that peak traffic occurs during an overlapping two hour interval for most users, and that about 1/2 of the total traffic on the network is sent during this time. Thus, during the peak hour of MINET operation,

we expect 1/4 of the total traffic to be sent through the network (1/2 of traffic divided by 2 hours in peak period). Using this information, we can convert the data in Figure 2-3 to yield peak traffic figures in msg/sec during the busy hour.

For each EM message that is sent by the user, a number of bits are transmitted from an EM host to a user (prompting, confirmation of delivery), and from the user to the EM host (entering addressees, text, editing text). Experience with EM on the ARPANET indicates that it typically requires 109K bits to be transmitted to the host and 164K bits to be transmitted from a host in order to compose a message and read 3 messages. This includes all control packets for subnet and host-host protocols (excluding internal network controls like routing).

There is also some traffic generated by the EM hosts on behalf of the users in the form of forwarded messages. We take the conservative assumption that half of the messages composed at a host must be forwarded to the other hosts (uniformly distributed over these other hosts). Intra EM host traffic is modelled as a message read (i.e., a message from EM host A to EM Host B generates traffic equivalent to reading the message from A). Applying these bit/message assumptions and the conversions for peak hour load further described in Appendix C, to the information in Figure 2-3 gives the peak data rates between sites that are shown in Figure 3-2. The figure represents traffic as a

Orig. Loc.	Contributed by	Nodal direction A	B	Util. of trunks on Canonical Path A to B	B to A
	% terminals				
LON	13.2	LON	LON	.079	.119
RDM	7.9	RDM	LON	.042	.064
BRM	11.4	BRM	FRK	.136	.206
FRK	10.5	FRK	FRK	.201	.304
RAM	12.3	RAM	LON	.122	.185
HDL	6.2	HDL	FRK	.052	.079
STT	3.5	STT	FRK	.037	.056
NAP	14.0	NAP	NAP	.137	.207
SIG	5.2	SIG	NAP	.082	.124
ROT	9.6	ROT	LON	.091	.138
IST	4.4	IST	NAP	.033	.050
ATH	1.8	ATH	NAP	.016	.024
	CONUS				
USA		USA	LON	.009	.014
USA		USA	FRK	.009	.014
USA		USA	NAP	.009	.014
	INTRA-EM				
LON EM		LON	FRK	.027	.031
NAP EM		LON	NAP	.024	.021
FRK EM		FRK	NAP	.030	.023

Figure 3-2 From/To Node Traffic Matrix (Busy Hour)

fraction of trunk utilization (capacity). The utilizations are shown for canonical paths between nodes over 9.6 Kbps trunks. The first set of entries corresponds to traffic contributed by MINET terminal users. The second corresponds to CONUS originated traffic homed to all three EM hosts. The third set of figures results from intra-EM host message forwarding.

The network configuration and topological design was developed next through an iterative process which tried to optimize the network with respect to the following goals:

- Keep the number of hops required to maintain a connection to the "homed" electronic mail host at a minimum by using a minimum spanning tree technique. The transmission delay on the low speed circuits is considered to be the most significant contribution to users' perception of service responsiveness. The network should not have any users more than two hops away from the homed host during normal conditions (assuming no hardware failures).

- Permit the network to grow through stages so that each new stage does not require changing previous stages' node placements, trunks, hosts, or subscriber homings. (All stages are purely additive.)

- Provide intra EM host and redundant path circuits.

- Keep the average loading on the trunk circuits (most critical overall resource) under .2 in order to be able to handle a substantial traffic growth.

- Provide at least two paths from every node in order to be insensitive to circuit failures.

The resultant network configuration is shown in Figure 3-3. This general configuration was chosen from 4 basic alternatives at the design workshop meeting.

3.1.3 Network Evaluation

The resultant network topology was extensively analyzed in order to create a final configuration and subscriber homing. The overall goals of responsiveness and reliability need to be more precisely defined.

Responsiveness is the amount of time a user at a terminal must wait for a system reply. The replies fall into three categories: (1) echoing characters typed on the keyboard, (2) host response to a command on a line of text entry, and (3) display of a message. In case 1, characters will be locally echoed by the TAC to which the terminal is directly connected. This is a full duplex path with minimal delay. The user will perceive immediate display of each character typed. Local

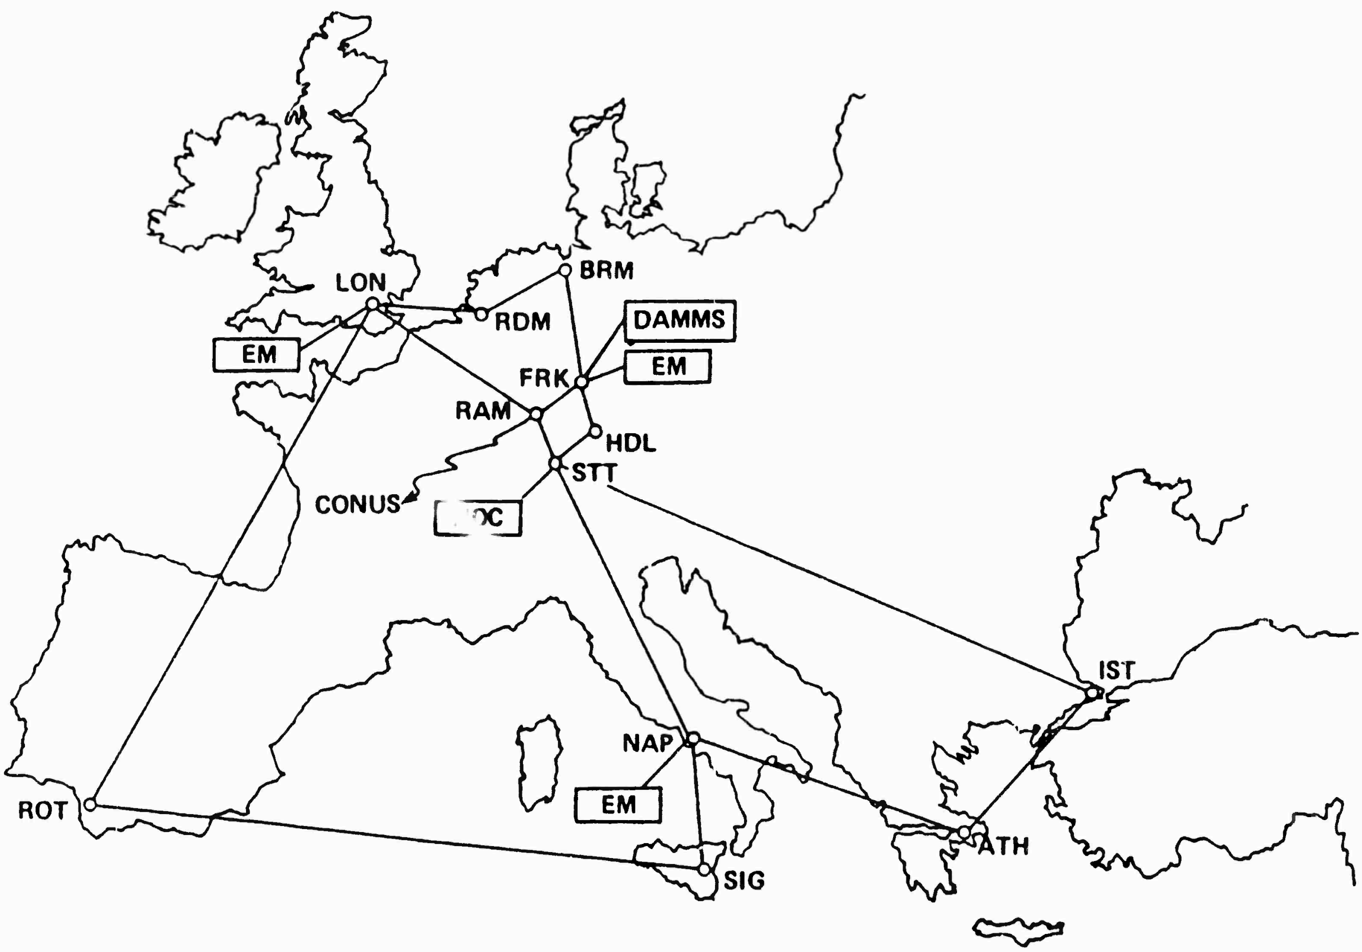

Figure 3-3 MINET Testbed Network Configuration

character echo is essential to the MINET design. All user data which passes across the network will be line-at-a-time transactions. In case 2, the network introduces delay for the line transaction. The user's line (command, text entry, etc.) will travel across the network, be processed by the host, and a line of text (reply, edit change, etc.) will travel back to the user. In case 3, delay is exhibited when the user requests many lines of output (e.g., message display). The wait will break down to a wait for the first line (same as case 2) and wait for subsequent lines. Most of the subsequent lines will arrive faster than a typical terminal can print them. (Short term throughput to a destination TAC will be approximately 4900 bps while most terminals can print at only 300 bps.) Thus the user will typically not perceive inter line delay. Line at a time reply delay (case 2) is, therefore, our critical performance parameter.

Studies have shown that users have the best productivity when system responsiveness is maintained below one second. Modest productivity reduction occurs between 1 to 3 seconds. Significant productivity reductions occur above 3 seconds. Above 3 seconds the operator's mind starts to wander, so he tries to do two things at once (e.g., maintain dialog with the computer and a conversation at the same time). A good MINET goal, therefore, is to keep system response under 3 seconds.

We will apply this goal to the line reply case. It must be pointed out that this is a narrow and strict application of the performance goal. The overwhelming majority of all system interactions will be quite responsive (cases 1 and 3). The components of the model case 2 reply are: user transmits a short packet, host processes a response, host transmits a medium (single line) packet. The host processing time can be assumed to be in the range of .25 to 1.0 seconds.

The design to performance goal, therefore, is that the model line reply test case round trip delay be under 2 seconds for a large percentage of terminals.

The reliability goal specifies the percentage of terminals which are not receiving the prescribed service under failure conditions. It is impossible to perform typical MTBF and MTTR analysis on the overall system since there is no equivalent network operating under similar conditions. We do not know the failure rates of circuits which we will use, nor the time it will take to get components fixed. The analysis must, therefore, simply predict performance when components are unavailable without attempting to predict component availability.

The specific reliability goal is that the performance criteria be met for 90% of the communicating terminals under any single component failure. Failures which stop service altogether

should be of a very low probability and should not affect any more than 15% of the terminals for each occurrence.

3.1.3.1 Performance Evaluation

We have estimated the network round trip delay for the test case scenario as seen by a user working on a terminal for each node under various traffic conditions. This is an extremely difficult calculation since there are many different, important, yet interrelated factors which can affect network transmission and propagation delay. A number of conservative assumptions were made in order to simplify things enough to make the analysis tractable. The details of the delay analysis are presented in Appendix D.

Figure 3-4 shows the cumulative percentage of terminals experiencing transmission delays less than or equal to the indicated values. It should be emphasized that these transmission delay calculations are intended to be representative rather than absolute, because in part they are based on our traffic estimates. The actual traffic patterns that eventually develop on the MINET will surely be different from those originally reported to us. Likewise, many other variables will also affect the individual delays experienced by any user of the MINET. We can, however, be reasonably confident that the maximum amount of round trip transmission-induced delay

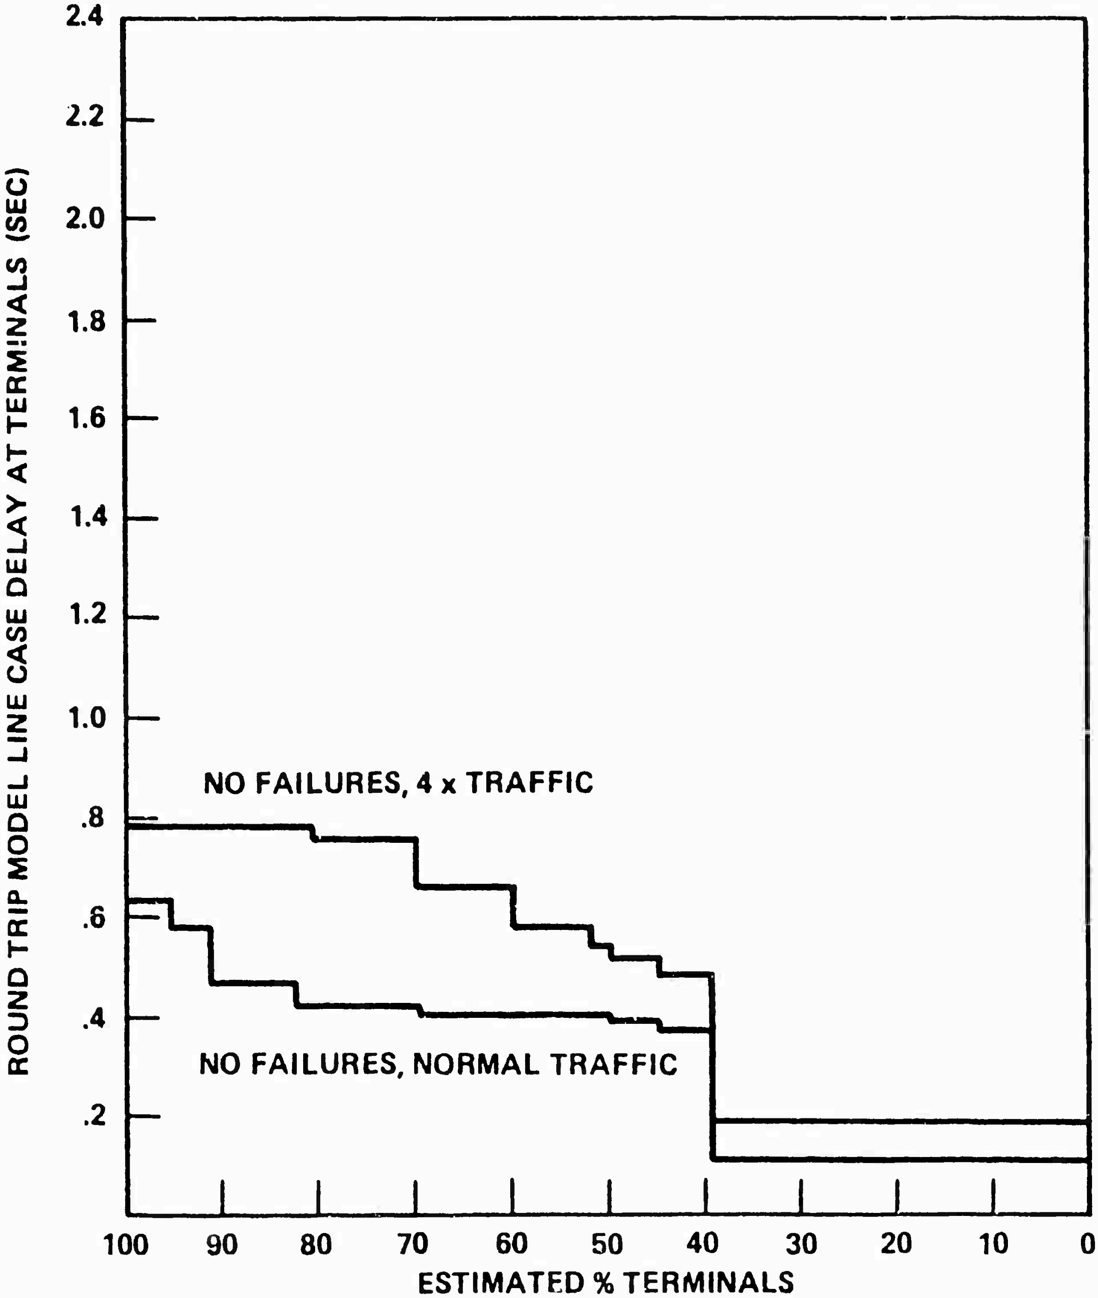

Figure 3-4 Estimate of MINET Performance
over Projected Testbed Period

should be less than or equal to the value we have predicted.

The traffic survey gathered data based upon the way the user community operates today. We can expect significant traffic growth during the life of the testbed network. The network will receive a suppressed service demand which currently exists. The users will probably find new uses for the network which cannot be predicted today. Finally, the network will undergo severe stress during periods of exercise. These considerations, taken together, yield a projected fourfold increase in the model traffic. Consequently, our performance evaluation also includes the case of a fourfold traffic increase.

Figure 3-4 shows the envelope of expected performance over the life of the testbed. The figure shows that the network responsiveness goal of 2 seconds can be maintained for all terminals, even with the fourfold traffic load increase.

3.1.3.2 Reliability Evaluation

Packet switched networks provide a very high degree of reliability under conditions of component failures. It was concluded during the design process that peacetime conditions generally result in a modest component failure rate. The reliability evaluation is therefore performed for all single component failures in order of their probability of occurrence.

3.1.3.2.1 Circuit Failures

Circuits are considered the most unreliable components in the system. The access circuits generally provide service to only a single terminal. Loss of an access circuit would, therefore, have minimal impact to the network or the subscriber community. Several nodes provide the capability to dial into a TAC. Critical users can be provided with dial-up terminals which can provide service during access circuit failures.

The network will automatically re-route traffic around a failed trunk circuit within seconds of the failure with no interruption of service to the users. All nodes except the gateway node in CONUS (#13) have at least two trunks connecting them to the rest of the network. The average connectivity (number of routes out of a node) is 2.5. The network will maintain full connectivity during any single and many double trunk failures.

Trunk failures which cause substantial re-routing of traffic, however, increase the hop count to a host and the utilization of the remaining trunks. Thus, although connectivity is maintained, responsiveness delay increases under certain failures. Fortunately, most circuit failures have a negligible effect upon performance. When critical circuits fail they affect a modest number of terminal users. Analysis shows that worst

case failures increase round trip test case delays to approximately 1.5 seconds for no more than 10% of the terminal population under normal traffic conditions.

The performance goal of providing better than 2 seconds responsiveness for a large percentage of the user population is maintained under all trunk circuit failures. This performance goal cannot, however, be met under the fourfold traffic load in the case of any one of approximately 5 critical trunk failures. It was not the intention of the design, however, that both high reliability and responsiveness be met under the fourfold traffic scenario.

3.1.3.2.2 Electronic Mail Host Failures

The second most probable system component failure can be attributed to the electronic mail hosts.

Each EM host will maintain accounts and mailboxes for a primary and secondary group of users. Therefore, each user will be "homed" to a primary host and a predetermined secondary host in case of primary host failure. The hosts are sized to provide such back up service to all of their subscribers. The system can, therefore, survive any single or even double host (after stage 2) failure with only modest inconvenience to the users. The users would have to log into their secondary host upon failure and then transfer their temporarily processed mail back to the primary host upon its recovery.

An EM host failure will substantially redirect traffic throughout the network. It will also add a significant percentage of traffic to the network from the users connected to the TAC which is collocated with the failed host. This traffic is normally kept off the network since it is local to a single node.

Detailed analysis shows that any EM host could handle not only its own traffic but also the traffic from the other two EM hosts. However, the network would not meet its performance goal because of the extremely high concentration of traffic which would exist on the few trunks servicing the remaining EM host.

3.1.3.2.3 Node Failures

The nodes are the most reliable component of the system. They are designed to operate completely unattended. They have no moving memory, terminals, or other typically unreliable components. Node failures can be divided into 4 types:

Type 1- Most node failures will typically affect only their local users. Under these conditions the affected users will completely lose service. It is easy to see from Figure 3-2 that the percentage of users affected is never over 15% for any single node failure.

Type 2- A node failure may also cause one EM host to become unavailable. This condition is the same as a single host failure already analyzed.

Type 3- Some node failures affect trunk circuits. Such failures behave the same way as circuit failures which were analyzed above.

Type 4- There are some low probability single component node failures which produce multiple trunk failures. These failures are considered so improbable relative to the other failures (i.e., circuit and host failures) that they have not been analyzed for delay performance impact at this time. However, even a worst case event of this sort, which is equivalent to the simultaneous failure of all trunks connected to a given node, will not partition the network, nor will it deny service to any user other than the users local to the failed node.

3.1.3.2.4 Overall Single Failure Performance

All single circuit and host failures are analyzed using the approach in Appendix D, with appropriately modified routes and trunk utilizations. This analysis is carried out only for the case of normal traffic. The result is shown in Figure 3-5, which is a revised version of Figure 3-4. The upper line shows the

composite of all worst case delays found. It, therefore, represents a region of expected performance under any single component EM host, trunk, and node (excluding type 4) failure.

It must be pointed out that there are many failure modes in which the performance is much better than the worst case. The worst case peak is produced by a single circuit failure. We cannot, however, predict where within the region of Figure 3-5 the network will operate. That critical circuit (Rota-London) may in fact be troublesome, although it is unlikely that it would be constantly failing. If that is the case, then the network should probably be re-engineered by adding another redundant trunk circuit.

The worst case upper limit of delay under single failures is still comfortably under 2 seconds for 90% of the terminals. The design is, therefore, shown to meet the conservative design specification both for responsiveness under fourfold traffic with no failures, and also for reliability under normal traffic with all probable single component failures.

3.2 Accessing Database Hosts

This subsection presents a strategy for connecting the four movements data base hosts (DB hosts) identified in the SOR to any ARPANET-like network, which includes the MINET. Specific issues that arise in connecting each of the four data base hosts --

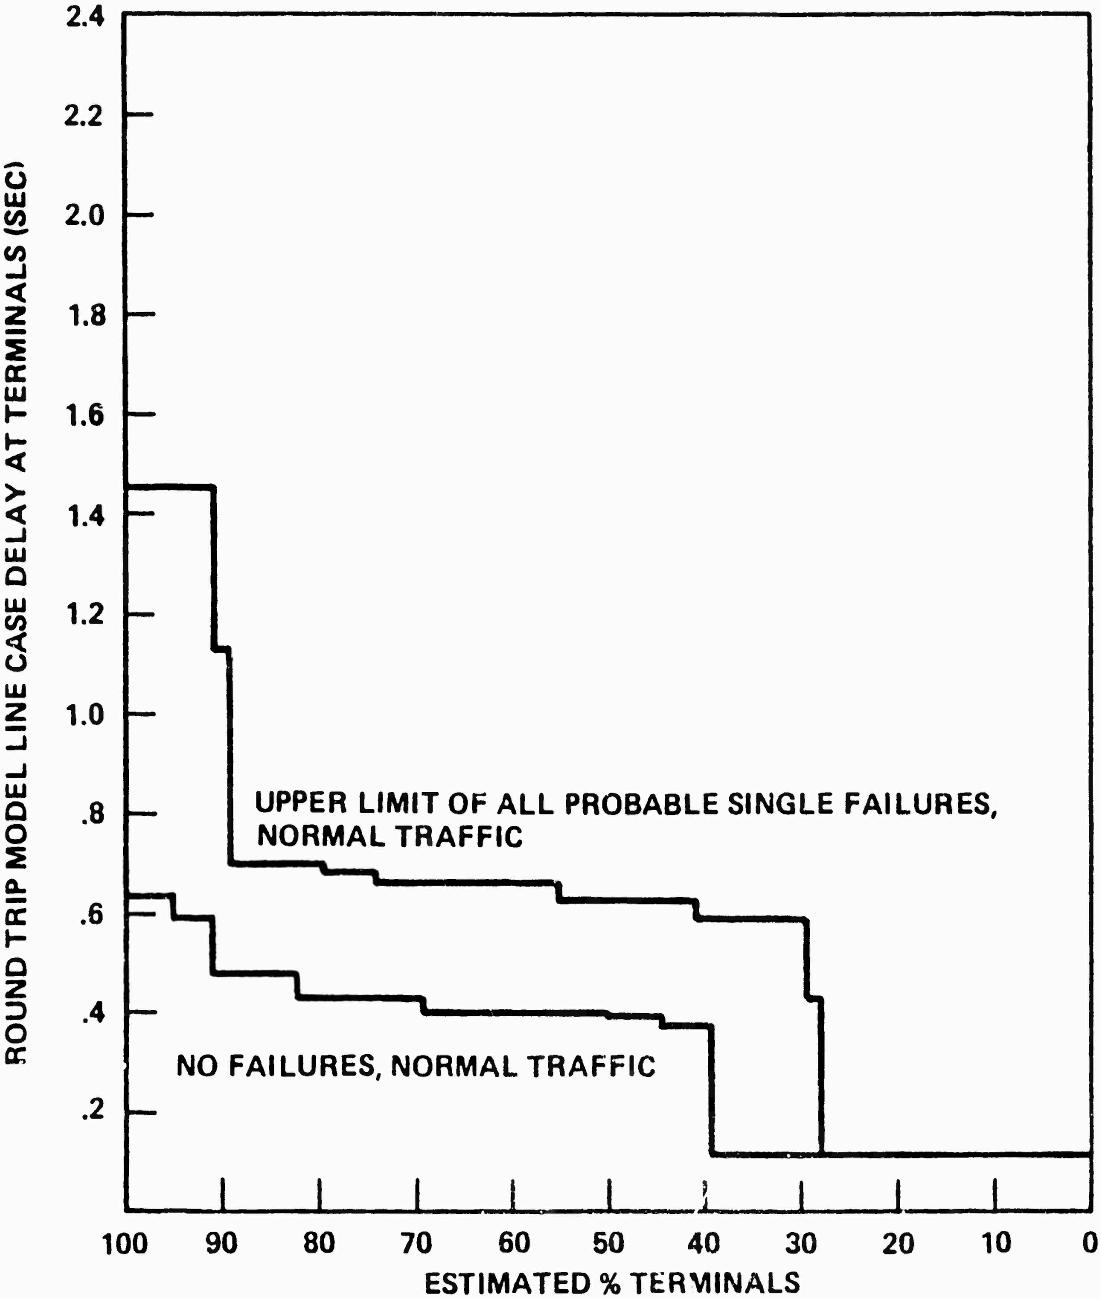

Figure 3-5 Estimate of MINET Performance over Range of Probable Single Component Failures

DAMMS, CAPS, TERMS, and the LIF, are covered in this subsection. However, a complete solution for interfacing any of these DB hosts is not included here. After the testbed MINET implementation contract is awarded, the contractor and USEUCOM can begin the detailed planning for attaching each of the four host systems, drawing on the general procedure included in this section. The Project Manager for the testbed MINET will sign Memoranda of Understanding with the Commands responsible for each of the DB hosts; these memoranda will specify the organizations that will do the interface work and the organizations providing the funding for the respective hosts.

The general concepts for connecting hosts to ARPANET-like networks are based on experience over the years with the ARPANET. The specific information about the four DB hosts is based on site surveys or telephone interviews. Both the DAMMS and CAPS sites were visited as part of the design study, and representatives at the LIF and TERMS sites were interviewed by telephone.

3.2.1 Issues in Connecting Host Systems

The effort involved in connecting a host system depends on the role of the host system in the network, as well as the current state of its hardware, its front-end processors, and its application software, teleprocessing software, and operating system. It also depends on the prospects for extending or

replacing the hardware or software in the future. If possible, one should attempt to capitalize on hardware and software development that has already taken place to connect similar hosts to the ARPANET. The cost of the hardware and software interfaces for an ARPANET connection may be spread over a number of host systems, if those hosts adhere to some standard. It is more likely that a potential network site can rely on earlier hardware and software efforts if that site consists of off-the-shelf hardware and software.

In the case of testbed MINET, two basic methods for connecting data base hosts are plausible. During phase 1, the electronic mail phase, connections to data base hosts are not necessary. Of course, during this phase, methods for connecting the hosts should be under active development. In phase 2, the query phase, it is possible to provide non-interactive queries to batch-oriented data base host systems by interposing the MINET EM hosts, as query and response spooling services. In phase 3, the update phase, the phase 2 approach can still be used. By phase 4, the file transfer phase, a full network connection is preferable to the EM host intermediary, since the mail format is not well-suited to arbitrary bulk transfers. The following subsections explore the specific issues that arise in connecting each of the four DB hosts to an ARPANET-like network.

3.2.2 The DAMMS Host

An IBM 4331, scheduled for installation in the summer of 1981, will serve as the host for the DAMMS (Department of the Army Movement Management System) application. The 4331 is an interim mainframe, however, and is expected to be replaced by an IBM 370/138 late in fiscal year 1982. The operating system for the 4331 will be either OS/VS1 or OS/VS2, running under VM370. Currently, all DAMMS input/output is local by means of either magnetic tape, punched card, or terminal access from IBM 3277 terminals.

Six methods for attaching DAMMS to the testbed MINET are suggested in this section. Following this, one of the methods is recommended as the near-term interface method for DAMMS. The six methods are:

1. an off-line batch method in which a collocated EM host maps EM messages into DAMMS batch-style updates and DAMMS batch-style output into EM messages. This method requires the support of an operator to move bulk media, such as magnetic tape, between the DAMMS and EM host systems.

2. an on-line batch approach in which a collocated EM host appears as an RJE (Remote Job Entry) terminal to the DAMMS host. This and the remaining approaches do not require any significant operator support.

3. an on-line batch method using a FEP and RJE terminals. This method uses a smaller dedicated FEP, rather than the larger EM host system, but it requires that some storage capacity and forms-handling capability be resident in RJE terminals, which must be placed at the various user locations.

4. an on-line batch scheme where DAMMS really is a testbed MINET host running the Internet Protocol (IP) and the Transmission Control Protocol (TCP), but where input and output spooling is still preserved and where RJE terminals are used.

5. an on-line interactive approach using a FEP and simple terminals. The application must be interactive in this case, and must use teleprocessing software that supports interactive terminals.

6. an on-line interactive scheme where DAMMS is a TCP host on the MINET. Here the application is also interactive, and relies on the TCP and IP protocols.

Of these six choices, the second choice appears to offer the lowest risk way of attaching the DAMMS host to the testbed MINET in the near term. The advantages of method two include:

- a minimum impact on the DAMMS host software. It will not be necessary to implement any of the ARPANET protocol software within DAMMS.

- provision of a well human-engineered single interface to the users, which can serve as a front end not only to DAMMS, but to other DB hosts. This interface includes the forms capability of the EM hosts.

- insulation of users from DB host availability and network connection issues.

- keeping high responsiveness. By design, the EM hosts are highly interactive, and local (not more than 2 intervening IMPs). The DB hosts, on the other hand, can be further away from the user. This consideration is especially important when connecting to the CONUS DB hosts, which are separated by many intervening IMPs, and a high-delay satellite network.

Perhaps the major drawback to method 2 is that it incorporates the delays inherent in the batch-style application program. However, the drawback is not too severe: the user may have to wait a few minutes before receiving the electronic mail message that contains the response to his query or update, but this service is a great improvement over the current methods. Electing method 2 at this time also does not preclude moving to a fully interactive approach, such as approaches 5 or 6, in the future.

3.2.3 The CAPS Host

Two systems, both Honeywell 6060s, serve as hosts for the CAPS (Consolidated Aerial Port Subsystems) application. One host normally processes cargo data, and the other processes passenger data. The operating systems for both machines are WWMCCS releases, which in turn are based on Honeywell's GCOS operating system. By fall of 1980 both hosts expect to be running WWMCCS system release 7.2. Terminals to each host are connected to Datanet 355 front-end processors. Plans call for replacing the Datanet 355s with Datanet 6600s. Both the cargo and passenger systems have local Honeywell VIP terminals connected to their front-end processors. There is also a remote connection from the cargo host to Dover AFB, a remote connection from the passenger host to McGuire AFB, and a shared remote connection to the Rhein/Main MAC site in Germany. All of these connections are 9.6 kilobit/second dedicated circuits. For the shared line to Rhein/Main, a Honeywell 716 processor acts as a switch between the two mainframes. A terminal line from one of the Datanet FEPs of each mainframe connects to the 716. Incoming requests from Rhein/Main are routed to the correct host according to addressing information in the protocol. Both the cargo and passenger hosts are scheduled to be connected to Autodin II.

A relatively low-risk method for attaching the CAPS hosts to the ARPANET or an ARPANET-like network is to use a network

front-end computer developed by Digital Technolog nc. (DTI) of Champaign, Illinois. DTI has implemented IP, TCP, and THP software under the UNIX operating system for a PDP 11/70 front end to H6000 series mainframes. (The THP protocol, developed for Autodin II, is similar to ARPANET Telnet protocol.) The PDP 11/70 front end, in turn, has an ARPANET 1822 interface. There is a host-to-FEP protocol between the H6000 mainframe and DTI's FEP. The software for this protocol has been written by Computer Sciences Corporation (CSC . This software is intended for use on WWMCCS hosts, and versions of this software are expected to track the future WWMCCS software releases. The protocol is sufficiently rich to allow the local host to open connections to remote hosts, and also to allow file transfers. Both DTI and CSC are working on projects which, as side effects, may improve the performance of this particular FEP approach. To use this approach, it will be necessary to replace the THP software with ARPANET Telnet software. Versions of the Telnet software for UNIX and TCP exist, such as the Telnet on the BBN-UNIX ARPANET host, which can be adapted for the DTI front-end processor.

3.2.4 The TERMS Host

Of the four data base host systems planned to be connected to the MINET testbed, the TERMS (Terminal Management System) host is the only one that is not yet operational. The planned host for TERMS is a Honeywell Level 6. Little information is

available about the TERMS configuration and software at present. One or more of the general methods for connecting the DAMMS host to the testbed MINET is likely to work for the TERMS host.

3.2.5 The LIF Host

The LIF (Logistics Information File) host system is an IBM 370/158 running IBM's SVS operating system. Plans call for upgrading to IBM's MVS operating system in 1981. The LIF host is not connected directly to the ARPANET, but it does have an indirect connection through the "OFFICE-1" ARPANET host in Cupertino, California. The connection is over a leased telephone circuit. The OFFICE-1 host is a PDP-10 running the TENEX operating system.

A user submits a query to the LIF data base by sending an EM message to a particular mailbox on the OFFICE-1 host. Some time later, the response to the query will appear in the mailbox of the user. Queries in the OFFICE-1 mailbox are transmitted periodically to the LIF host by means of an ad hoc polling protocol. The queries are processed, and the responses are returned, on a periodic basis, to the OFFICE-1 host. The OFFICE-1 mail software sends the responses from the LIF host to the users who submitted the corresponding queries. It is the responsibility of the user to know the correct format of a query against the LIF data base. This method for connecting the LIF

host to the ARPANET is similar to the recommended method for connecting the DAMMS host to the MINET.

Thus the LIF host is, in this indirect fashion, able to be queried at the present time. The LIF uses another system as a front end. This FEP strategy for the ARPANET connection can also be used to connect the LIF to any ARPANET-like network. However, the query-response service is not interactive; responses may take up to 40 minutes or more to be delivered. In addition, the requirement that data sent to the LIF be in the form of EM messages may restrict the kinds of file transfers that can take place between the LIF and other data base hosts in phase 4 of the MINET testbed. Thus, a direct network connection may be needed to meet the long-term requirements of the MINET user community. The alternatives suggested for connecting the DAMMS host are also applicable to the LIF.

3.3 Options for the Electronic Mail Service

The selection of an appropriate electronic mail service for the testbed MINET is an extremely important task. Its acceptance by users and its performance will weigh heavily in the evaluation of the usefulness and success of the testbed MINET. This section discusses the selections we made, the candidate electronic mail host systems and electronic mail services we evaluated, and the specific criteria we used in evaluating them.

The crucial nature of the electronic mail service influenced the selection criteria that we established: we wished to employ an electronic mail service whose overall functionality and style of user interaction took advantage of the body of experience in the ARPANET mail community; we wished to employ a host computer system for the electronic mail service that has been used extensively with the ARPANET; and we wished to employ an electronic mail service that persons without previous computer experience would find relatively easy to learn and to use. Further, we wished to employ an electronic mail service whose interaction style would result in network traffic between the user's terminal and the electronic mail host that would make relatively efficient use of the MINET facilities. In addition, we wished to employ an electronic mail service that was efficient and well-matched to the host computer system we selected, in order to provide reasonable performance. Finally, we wished to select an electronic mail service that supports those features requested by the respondents to the functional requirements survey. Electronic mail services that do not meet the functional requirements presently, but which can be modified to meet them without undue effort, are also worthy of consideration.

These selection criteria argued against the development of new electronic mail software, against the use of existing electronic mail services with complex and perhaps difficult user

interaction styles, and in favor of the use of a host computer system widely accepted in the ARPANET community paired with an electronic mail service that was well suited to ' .

3.3.1 Selecting the Electronic Mail Host

In today's ARPANET environment, electronic mail services are most often found on general-purpose timeshared server hosts, as one of a broad range of services provided. These server hosts range in size -- and cost -- from "high-end" minicomputers such as the Digital PDP-11/70 and VAX-11/780 to large-scale machines such as the Digital DECsystem-20 and Honeywell 68/80 Multics. The amount of electronic mail traffic that each of these server hosts can support depends upon the aggregate load placed upon the machine by the users of all its services, and is, in general, not limited by the characteristics of the electronic mail service itself.

In the MINET, on the other hand, the server host providing the electronic mail service will be performing only that function, so a smaller server host could support a good deal more electronic mail traffic than is usual in the ARPANET environment. Analysis of user traffic requirements indicated that the initial electronic mail needs of the MINET would be served adequately by a small number of high-end minicomputers such as those mentioned above. The high-end minicomputer is also an economical choice,

as will be shown in section 4 of this report. As use of the MINET electronic mail service grows, additional high-end minicomputer hosts can be added incrementally, at reasonable cost. Finally, electronic mail service provided by two or more smaller hosts will have greater overall reliability than service provided on a single large-scale host.

The time-shared operating system for high-end minicomputers that enjoys the widest use and acceptance by users in the ARPANET community today is UNIX, developed by Bell Laboratories and available through Western Electric. We have chosen UNIX as the operating system for the MINET electronic mail hosts for several reasons. UNIX presently runs on a variety of computers from several vendors, including several models of the Digital Equipment Corporation PDP-11 family and VAX-11/780, as well as the BBN Computer C/70. Hardware is readily available to connect all of these machines to ARPANET-like networks. Several electronic mail services have been written for UNIX and are in use in the ARPANET community today. We believe that these facts, taken together, make UNIX a compelling choice for the operating system of the MINET electronic mail hosts.

3.3.2 Requirements for the Electronic Mail Service

Our next task has been to select an electronic mail service for the UNIX electronic mail hosts. Our primary objective has

been to select a software system that can provide an efficient service, easily used by personnel without previous experience with computer terminals and computer-based systems. The selection criteria we have used fall into three categories: (1) functionality and features, (2) implementation efficiency, and (3) availability and support.

In the interest of saving both development time and funds, we decided against proposing the development of electronic mail software especially for the MINET. Although such an effort would result in a software system that would make efficient use of the MINET communications facilities and provide exactly the functionality and user interaction style needed for the MINET, it would be a time-consuming and expensive task. Instead, we examined three available electronic mail services for UNIX: INFOMAIL, a commercial product of the BBN Information Management Corporation, MH (for "Message Handler"), developed by the Rand Corporation under a Defense Department contract, and MSG/SNDMSG, UNIX versions of two older, popular ARPANET mail-handling programs, implemented for UNIX by BBN.

In the area of functionality and features, we asked the following questions:

1. Is the system designed for the majority of MINET users; that is, people who have had no experience in the use of computer-based systems and computer terminals?

2. Does it provide:

   - a simple message editing capability?
   - the ability to construct forms and templates, or use forms and templates constructed by others?
   - "bulletin boards"; that is, collections of messages "posted" for examination by groups of users?

In the area of implementation efficiency, we asked:

1. Was the service designed specifically for UNIX, or is it a "portable" system adapted to UNIX?

2. Was it written in a language appropriate to UNIX?

3. How does it organize message storage; does it take advantage of the UNIX file system and directory structure?

4. Does it store a single copy of each message, or does it store a separate copy for each recipient?

5. Is its interaction style appropriate for the MINET; that is, is the network traffic generated by users interacting with the service on the electronic mail host of a type that is handled efficiently by the MINET communications subnetwork?

Under availability, we asked these questions:

1. Is the service available now? If not, when will it be available?

2. What does it cost to procure?

3. What organization currently maintains the software? Is continued support expected, or will support need to be provided specifically for the MINET?

3.3.3 Electronic Mail Service Evaluation

We now present our evaluation of each of the three electronic mail services we considered, followed by our recommendations.

3.3.3.1 INFOMAIL

INFOMAIL is a forthcoming product of the BBN Information Management Corporation. It is intended to be a portable software system that can be run on a several different types of host computers. It is written in FORTRAN, a language which INFOMAIL's implementors feel is the most widely available language in which it is feasible to write a software system as complex as INFOMAIL.

INFOMAIL builds its own message storage structure on top of the host system's file system. In this structure, it stores only a single copy of each message posted in the system, regardless of

the number of recipients of the message. Each user's "mailbox" is then a data structure containing pointers to the single system-wide copy of each of the messages that one would normally think of as being "in" the mailbox. Because of its message storage system, INFOMAIL should be economical in its use of storage space. However, its performance and overall efficiency will depend a great deal on the efficiency of its interaction with the UNIX file system.

INFOMAIL is designed specifically for use by non-computer users -- the MINET type of user. Its designers feel they have taken special care in the design of the user interface to provide what might be termed a "friendly" user interface. For example, the user can re-define both command verbs and the names of the various fields of messages. This ability could be used in the MINET to change INFOMAIL command verbs and message field names to use the terminology of the logistics community. For example, "To:" and "cc:", as one might find them in the commercial world, could be re-defined to be "Action:" and "Info:". INFOMAIL uses text editing facilities provided by the host to aid the user in composing messages. UNIX provides several text editors; an easy-to-use editor could be selected for use with INFOMAIL.

INFOMAIL does not provide a "bulletin board" capability; multiple-recipient distribution lists must be used instead. However, because of INFOMAIL's message storage system,

distributing the same message to a large number of users does not result in replication of many copies of the message, so the lack of a bulletin board facility does not engender a storage cost penalty.

INFOMAIL provides forms and templates for easy creation of standardized messages. There is a standard procedure for creating new forms and making them available to the user community.

INFOMAIL is capable of supporting the line-at-a-time interaction style of the MINET terminals. It does not require the less efficient character-at-a-time interaction style.

INFOMAIL will be available for UNIX in the first quarter of 1981, which is certainly within the MINET timeframe. It is a software product of the BBN Information Management Corporation, which will provide full support for it. As a software product, both the software system and a yearly maintenance and upgrade contract must be purchased from the vendor.

3.3.3.2 MH

MH, which stands for Message Handler, was developed by the Rand Corporation under a Defense Department contract, and is described in the Rand Corporation Report R-2367-AF of November 1979. MH was developed by Rand specifically to provide an

efficient electronic mail service for UNIX. Development of MH was motivated by the poor performance and consequential poor user acceptance of a predecessor system, MS ("Message System"), that was developed to explore user interface issues in electronic mail services. MH has enjoyed excellent user acceptance at Rand and is currently being made available by Rand to other ARPANET UNIX sites.

Since MH was designed with efficiency on UNIX in mind, it uses, internally, a number of UNIX system features. In particular, its user interaction style is like that of the UNIX system itself (the UNIX "shell"), and it adopts the directory tree structure of UNIX for its message storage structure. The user interaction style of MH is therefore tailored to the UNIX user, rather than to the MINET-type user. However, the versatility of the UNIX shell command language makes it possible to evolve a user interface for a MINET MH that would be appropriate for the MINET user. MH itself is written in C, the implementation language of choice for UNIX.

Like INFOMAIL, MH allows its users to edit messages using any available text editor. An appropriate editor would be selected for MINET from those available. Like INFOMAIL, MH has extensive capabilities for forms and templates, a feature important for the MINET. In addition, MH provides a well-developed bulletin board facility. MH does support the

line-at-a-time interaction style planned for the testbed MINET.

MH stores a message copy for each message recipient. For example, a message sent to three users on the same host would be stored as separate copies in each of three distinct mailboxes. This message-per-recipient implementation can consume large amounts of secondary storage if broadcasting of EM messages is to be commonplace in the MINET.

MH is available at the present time from Rand Corporation; Rand will supply all of the MH software on a magnetic tape. The only cost to the U.S. Government for MH is a small fee to help cover the distribution cost. If MH is selected for the MINET, it will be necessary for the contractor to support it. The manpower estimate for the testbed MINET program plan provides for a programmer in the theater who will be responsible for EM software support.

3.3.3.3 MSG/SNDMSG

The MSG/SNDMSG programs exist, in slightly different forms, on a number of ARPANET hosts. UNIX-based versions of MSG/SNDMSG are available from existing ARPANET sites, at only a very slight cost to the Government. The implementation of MSG/SNDMSG discussed here is the one that exists on the BBN UNIX system. Unlike the INFOMAIL and MH user interfaces, the MSG/SNDMSG user interface cannot be modified easily to meet a particular set of

user requirements. On the other hand, the command interface is not oriented towards any specific computer system.

MSG/SNDMSG allows the user to invoke a text editor of his choice when composing and forwarding messages. It supports distribution lists for message broadcasting. It does not provide forms and templates, nor does it provide a "bulletin board" capability. MSG/SNDMSG for UNIX can operate in the line-at-a-time mode of the MINET testbed.

There are versions of MSG/SNDMSG written specifically for UNIX, in the "C" language. Like INFOMAIL, MSG/SNDMSG organizes its message storage on top of the UNIX file system. Like MH, MSG/SNDMSG stores a copy of the message per recipient.

The MSG/SNDMSG software from the BBN UNIX site is available to the Government for only a small distribution cost. This software is currently being supported at BBN and is likely to be supported throughout the period of the testbed MINET program. However, in-theater programming support from the contractor will be required in order to extend the user interface to meet the requirements of the testbed MINET users.

3.3.3.4 Recommendations

Both MH and INFOMAIL provide a more flexible user interface than MSG. In the MINET testbed, the ability to adjust the user

interface, in the field and without extensive effort, will be important. INFOMAIL has advantages over MH; it will be a supported commercial product, and it stores a single copy of each message. On the other hand, MH is available today and INFOMAIL is still under development. Thus the recommendation for electronic mail software is the best off-the-shelf choice, which is MH. Once the development of the testbed MINET begins, INFOMAIL should be reconsidered to assess whether it has become a stable product.

The budgetary cost estimates in thi[illegible]rt are based on th[illegible] assumption that the mail service chosen is [illegible] This assumption is a conservative one, since choosing INFOMAIL would require less contractor manpower for development and support.

3.3.4 Interfacing Electronic Mail and Telex Services

Interfacing the electronic mail service of the testbed MINET to the Telex network will allow shippers, consignees, and other sites that do not have testbed MINET terminals to communicate with those sites that have terminals. The interfacing plan calls for terminating Telex network circuits at one or more of the EM hosts. Connecting a Telex circuit directly to a computer, in turn, will require PTT approval. If the direct Telex-EM host connection is not approved, some "off-line" methods for interfacing Telex and the testbed MINET are also possible,

although these will require manual intervention from the testbed MINET operators.

A second level of addressing, or "readdressing," is required by the Telex-MINET testbed interface plan. Incoming messages from the Telex network to the testbed MINET have a 2-level address, in which the outer address is a standard Telex number that specifies a particular EM host, and the inner address specifies the actual destination mailbox or address list, probably in just the same format that would be used for the mail service within the testbed MINET. Outgoing messages to the Telex network also have 2-level addresses, in which the outer address is an EM address to a special "Telex server" network mail destination, and the inner address is the actual Telex number of the destination party.

In addition to the UNIX operating system, the testbed MINET software and the EM software, some special-purpose software modules will be required to interconnect the EM and Telex services. Two major special-purpose modules are described here briefly.

First, there is a "Telex Readdressing program." This program deals with incoming messages from the Telex network, as well as with outgoing messages to the Telex network. For an incoming message, a receipt can be returned to the sender, if

requested. The inner address is examined, and a copy of the message is sent to each recipient. If the inner address on an incoming message cannot be determined, the message is given to the 'Librarian" program, described below, which will notify an operator. For an outgoing message, the Readdressing Program will extract the inner address, which is a Telex number, and place the message into a mailbox that corresponds to that Telex number. On a frequent basis, outgoing messages will be extracted from the mailboxes that correspond to Telex numbers, and will be sent to the Telex terminal over an auto-dial interface to the Telex network.

Second, there is the Librarian program, which maintains addressing tables and system parameters. In addition, it will receive all messages for which the Telex readdressing program could not find a destination, and will provide software tools to aid the operator in redirecting these messages.

Since incoming and outgoing messages are stored in mailboxes for a brief period of time, they will not be lost due to line failures. Storage for messages to and from the Telex network, however, is limited. The current design does not call for long-term archiving of these messages. The electronic mail system will keep track of the age of all messages -- including messages to and from the Telex network -- that are in mailboxes. Subject to some user override capabilities, messages with ages

older than an administratively set value will be deleted from mailboxes.

3.4 Internet Architecture

Three options for providing access to the CONUS data base hosts were considered in this study. The first option is to extend the ARPANET to Europe to connect to DAMMS and to the MINET terminal users. The second option is to build a single network, distinct from the ARPANET, that covers both Europe and CONUS, and that ties together all MINET hosts and users. The third option is to build a separate net in Europe and connect it to the ARPANET through internet gateways. The third option is the recommended one.

The first option achieves the objective of connecting the MINET community with the least amount of extra hardware and circuits. However, it should be dismissed on the grounds of management and control. The testbed MINET, and ultimately the MINET, must be responsive to the needs of the user community in the European theater, and in order to achieve this responsiveness, the facilities must be controlled locally. Attempting to support the needs of the diverse ARPANET user community and the testbed MINET user community will present very difficult administrative dilemmas, especially in exercise situations. The testbed MINET also requires a local NCC under

normal conditions in order to manage circuit failures and to contact repair personnel. The testbed MINET management must maintain jurisdictional control over the network resources.

The second option is more appealing since the network can be dedicated to the MINET user community. The major drawback to this approach is duplication of facilities. Trunk circuits and IMPs for the CONUS part of such a network would be used to provide connectivity that can already be provided by the ARPANET. One would lose the advantage of sharing the existing ARPANET facilities. The second option can allow the MINET to be a dedicated network, but at a higher cost.

The third option, internetting a European MINET to the ARPANET, is a good compromise between the first two options, especially for the MINET testbed. This option allows the entire European net to be dedicated to the movements activity, while avoiding duplicating ARPANET facilities in order to connect to the CONUS hosts. The three CONUS hosts would be brought onto the ARPANET, and be reached by MINET users through the MINET-ARPANET gateway. This choice does not preclude providing dedicated MINET facilities in CONUS at a later time, once the testbed network has been demonstrated to be effective.

Having selected the third option, it is possible to connect the MINET directly to the ARPANET, or indirectly through an

intermediate network such as SATNET. However, using intermediate networks reduces the degree of control that the MINET user community can have over its facilities, and makes the data protection requirement more difficult to satisfy. Thus it is recommended that the MINET connect directly to the ARPANET through an internet gateway.

At present, only one trans-Atlantic circuit has been identified to connect the MINET to the ARPANET. This circuit, which has been designated by DCA Europe to be for the MINET-ARPANET link, is a 50 Kbps circuit connecting Landstuhl, West Germany and Ft. Dietrick in the CONUS. Once the MINET testbed development is underway, a second MINET-ARPANET circuit should be identified and scheduled to be in operation when the first CONUS MINET host is available for phase 2 remote query.

Since the jurisdictional boundary between the ARPANET and the MINET is defined by the location of the gateway computer, the side of the Atlantic Ocean on which the gateway is installed will determine whether the trans-Atlantic link is part of the ARPANET or the MINET. It is recommended that the link be under the jurisdiction of the MINET, so the gateway should be in the CONUS. Collocated with the gateway will be a MINET IMP. Both the gateway and the IMP, logically parts of the MINET, can then be maintained in the CONUS. The other side of the gateway will be connected to an ARPANET IMP. This connection could be either a

very distant host link to an existing ARPANET IMP, or a local host connection to a new IMP installed together with the gateway. The new IMP can be connected into the ARPANET as a stub, or with two links for greater reliability. Obtaining this new IMP increases the cost of the program, but provides more complete reporting of the state of the intervening link by means of the line up/down protocol for IMP-IMP circuits.

Connecting the MINET and ARPANET with a gateway affects the choice of networking protocols. The DoD standard Internet Protocol (IP) and Transmission Control Protocol (TCP) will be used to allow addressing of host systems on other networks. The original ARPANET Network Control Program (NCP), still in wide use on the ARPANET, is adequate for communication only within a single network. A program is presently underway to make IP and TCP available on all ARPANET hosts to support internetworking. Phase 1 operation of the MINET could be supported by NCP, since no internetwork communication will take place. However, the following phases will require IP and TCP, so the usefulness of NCP would be short-lived. Accordingly, IP and TCP are to be provided in each of the MINET hosts. Implementation issues are discussed in section 3.2.

Higher layers of protocol must also deal with the internetwork environment. The Telnet protocol, which allows terminal access to hosts, the File Transfer Protocol (FTP), and

electronic mail protocols must all be able to address internet hosts. Both the EM hosts and the data base hosts will need to support Telnet and FTP. The internet version of the electronic mail will reside on the EM hosts. Presently, internet versions of Telnet and FTP exist for UNIX-based hosts, although this software will need some tailoring for the MINET user community. Internet mail addressing protocols are currently under development.

The recommended internet architecture can be expanded in the future to connect the MINET to AUTODIN II, either in the CONUS or in Europe. The gateway solutions being developed by DCA to interconnect the ARPANET and AUTODIN II will be directly applicable to the MINET, since it is an ARPANET-like network. All host protocols in the MINET are being chosen from the DoD standard protocols which support multiple network operation regardless of the networks involved. Thus the MINET can be connected to AUTODIN II by a gateway with minimal impact on the hosts and users.

3.5 Data Protection

As mentioned in the functional requirements section, only unclassified traffic will be sent through the MINET. Nevertheless, a minimum requirement is that testbed architecture protect against hostile exploitation of aggregate logistics

information transmitted through the network. Thus, at a minimum, the major links of the MINET testbed need to be protected in such a way as to prevent successful passive wiretapping attacks. In a passive attack, an intruder merely observes messages on the communication links. An active attack, on the other hand, is one in which an intruder can modify, delete, delay, reorder, replay, or fabricate messages on the communication links.

It should be emphasized that the requirements do not state that the MINET testbed circuits be protected against active wiretapping attacks. It is sufficient to protect against hostile exploitation of the information in the MINET. Thus, the objective for the testbed effort is to prevent successful passive wiretapping attacks, but not to protect against active attacks. However, the data protection strategy for the MINET testbed has been designed so that, if protection against active attacks is desired at a later time, the extension can be made with little effect on the MINET operation or terminal user interface.

The basic choices for protecting information in the MINET are end-to-end and link encryption techniques. End-to-end techniques protect the data from some point near the source host to some point near the destination host. In pilot projects sponsored by DARPA to gain experience with end-to-end methods, the encryption and decryption is provided by a unit between each host and its neighboring IMP. Link-oriented techniques protect

the data on the links between IMPs. This is accomplished by placing encryption/decryption units at each end of the IMP-IMP link, close to their respective IMPs.

The only end-to-end equipment that is likely to be off the shelf and approved for MINET use is the Private Line Interface (PLI) developed by BBN. These units, however, are not well-suited to the MINET application: they must be placed in secure locations, and they must be tended only by individuals with the appropriate high-level clearances. Furthermore, providing one of these units for each EM and data base host would practically double the hardware portion of the estimated testbed MINET budget.

On the other hand, there exists a wider choice of off-the-shelf link encryption equipment. The two most promising classes of equipment are cryptographic devices which are Government Furnished Equipment (GFE), and devices based on the U.S. Federal Data Encryption Standard (DES). Of these, the DES-based devices are more attractive since they can be used in an uncleared area, while the GFE cryptographic devices must be kept in a cleared facility, and tended only by cleared individuals. DES-based devices have received DoD approval only for unclassified traffic, but this limitation is acceptable for the unclassified MINET. The DES-based devices considered for the MINET also support remote key loading, which simplifies network

operation. Finally, the cost of the DES-based units is much less than that of the GFE units. For these reasons, DES-based devices are recommended as the means of providing data protection in the MINET.

The data protection plan is to equip each trunk circuit in the MINET with DES-based link encryption devices. There will be a DES unit at each end of the trunk circuit, between the modem and the IMP. The only exception to this plan is to equip the trans-Atlantic circuit to the CONUS with GFE cryptographic units rather than DES units. In the case of the 50 Kbps circuit to the CONUS, GFE units must be used since the currently available DES units can operate only at rates up to 9.6 Kbps. The DES units will encrypt all data above the link level. That is, the special characters that are used to provide framing for the message text pass through the unit unencrypted; the message text and its checksum are encrypted. To do this, the DES units must incorporate knowledge of the link level portion of the IMP-IMP protocol. Off-the-shelf units do not support the IMP-IMP link level, but they do support the BISYNC protocol, which is similar. For a small fixed charge, it is possible to modify the contents of the read-only memories in the units to accommodate the IMP-IMP link level protocol.

For terminals which must connect to their nearest TAC over telephone circuits, it is possible to provide link encryption

using a different version of the DES unit that encrypts data transmitted in asynchronous, 8-level character format. In this case, each character, together with its start and stop bits, is encrypted. DES units for asynchronous terminals are available off the shelf.

A potential problem area, when incorporating cryptographic units into the MINET, is the effect of data errors on the MINET circuits. Some kinds of units lose cryptographic synchronization in the face of data errors on circuits, and require manual re-synchronization. The DES units considered for the MINET employ self-synchronizing ciphers, and are thus able to recover from errors on the lines. The commercial DES units employ the cipher feedback (CFB) mode of the DES algorithm, which limits the propagation of errors so that if a single character is garbled, the decrypting unit falls out of crypto synchronization for the next eight characters. At that point, the encrypting and decrypting units will once again be in crypto synchronization. Therefore, the increase in packet error rate is negligible when using such DES units.

Another possible problem area is the interaction of loop-back testing of circuits with the cryptographic equipment. In order for loop-back testing to be done without manual intervention, the crypto unit must pass loop-back signals from the IMP to the modem. The MINET DES units, which will employ a

25-pin CCITT V.24 (RS-232C) electrical interface, should be able to pass signals on those pins of the 25-pin connector which correspond to the local and remote modem looping commands. Otherwise, some kind of break-out box to pass these signals around the crypto unit would be required. In some contexts, allowing such signals to pass through a cryptographic unit might violate red-black separation policies, which strictly limit information flow from the red (IMP) environment into the black (telephone circuit) environment. In the case of the unclassified MINET, ease of fault isolation is the overriding concern, making transparent passage of loop test signals through the crypto unit a very important feature. Units considered for the MINET provide this support.

The DES units for the MINET support down-line key loading. This is accomplished using a two-level keying scheme. The primary key changes only infrequently and is used to protect the secondary key. The secondary key changes more frequently -- perhaps once a day. To perform a secondary key change, one DES unit assumes the role of master and the other the role of a slave. The master unit initiates the secondary key change sequence, using a special protocol. All exchanges in this sequence are encrypted under the primary key. An operator is required to activate the key loading sequence at the master unit. In order to simplify operations, certain MINET nodes can be

designated as nodes with master crypto units for purposes of key loading. The master nodes should be selected in such a way that every node in the network is either a master node itself, or else between two master nodes. This key loading scheme eliminates the need for physical transport of secondary keys between any of the testbed MINET nodes. In the 12-node MINET testbed, for example, master crypto units at London, Bremerhaven, Frankfurt, Stuttgart, Sigonella, and Athens could be used to down-line load secondary keys into the remaining slave units.

4. PROGRAM PLAN

The preceding sections of this report showed first how to take the requirements of the MINET users and map them into requirements for the network itself, and secondly, how to take the requirements for the network and map them into a network design that satisfies those requirements. In this section, we suggest how to obtain the necessary circuits and hardware for the network design, how to operate the network, and we present a project plan for the testbed development, installation, and operation.

4.1 Circuits

Acquisition of the required number of network trunk circuits of adequate quality is essential to the success of the MINET testbed. DCA Europe will serve as the network control service for the MINET, which will include the responsibility for both acquisition and the maintenance of the circuits. As indicated previously, there will be a mix of military (DCS) and commercial (PTT) circuits. As soon as funding is available for the MINET testbed implementation, a contact within DCA EUROPE should be identified as the MINET circuit manager. The first task of the MINET circuit manager will be to process the requests for the stage 1 trunk circuits, which connect the IMPs, as well as for the stage 1 tail circuits, which will connect the TACs to remote

terminals. For each of the PTTs that will provide circuits for the MINET, there should be a single point of contact for acquisition and one for maintenance. The MINET circuit manager at DCA Europe will deal with the PTTs through these identified acquisition and maintenance contacts.

The trunk circuits, all of which will support a data rate of 9.6 Kbps full duplex, must be 4-wire circuits. The tail circuits, which will support data rates of 300 to 1200 bits/second full duplex, may be either 2-wire or 4-wire circuits.

It is recommended that all of the circuits used for trunks in the MINET testbed be M1020 conditioned circuits. This kind of conditioning has been required in the European theater in order to obtain a 9.6 Kbps data rate over voice-grade circuits. Manufacturers of some of the newer 9.6 Kbps modems on the market claim that their products can operate at reasonably low error rates in point-to-point mode even on unconditioned voice-grade circuits. Nevertheless, M1020 conditioning is recommended for the following reasons:

1. M1020 conditioning provides a specification of circuit quality. Without conditioning, there is no specification of quality; a circuit "works" if it is possible to carry on an intelligible voice conversation from one end to the other. Without conditioning, one cannot claim that circuit quality is substandard, since there is no standard.

2. The claims made by some manufacturers that their products can tolerate unconditioned lines may not hold up when dealing with the various DCS and PTT circuits that will make up the MINET trunks, even though such claims may be valid for the U.S. phone system.

3. In keeping with the low-risk policy of the MINET testbed, obtaining M1020 conditioned lines is a cost-effective investment. The additional charge by the PTTs for this service, at least in West Germany, is modest.

Trunk circuits, which are logically disjoint in the schematic representation of the MINET testbed, should be physically disjoint. The Rotterdam node, for example, has trunk circuits connecting it to London and to Bremerhaven. If these two circuits were routed through the same conduit once they were sufficiently close to Rotterdam, then damage to the conduit might cut off both trunk circuits. The best possible scheme for connecting multiple trunks to a given IMP is to have physically distinct circuits right up to the rack containing the modems. The next best approach is to keep the trunk circuits physically distinct at least until they are within the particular U.S. military base where the IMP is located. Circuit conduits within these bases should be easier to monitor and repair than those on the outside. It is important to include in the circuit request to DCA Europe that the trunks be physically disjoint. DCA Europe

will then attempt to provide as great a degree of physical separation as possible.

4.1.1 PTT Circuits

The recommended trunk circuit configuration for the MINET is to use half PTT and half DCS circuits. However, the likelihood of obtaining very many DCS circuits for the MINET testbed is low. Most of the circuits in the testbed will be leased from the PTTs. Once the trunk and tail circuit requests are presented to DCA Europe, DCA will carry out a preliminary routing investigation. DCA will attempt to use DCS circuits wherever possible, and submit the remaining circuit requests to the PTTs.

When requesting circuits from the PTT, DCA can specify a future date when it will need the service. Lead time for acquiring the PTT circuits is about 6 to 9 months, so the circuit requests should be submitted as early as possible in the MINET testbed program.

One of the major administrative issues that must be confronted in installing the MINET testbed is the notion of Host Nation Approval. Any device that is to be attached to a PTT circuit must be approved by the PTT of the host nation. Such an approval process is called "homologation," and a device which has received such approval is called "homologated." These homologation regulations apply in each of the nations where

testbed MINET circuits are planned. In general, each host nation has a somewhat different list of approved devices. Since many of the MINET trunk circuits will cross international boundaries, it is necessary on those circuits to use devices that are approved in each of the end-point countries. For the MINET, the host nation approval policy affects the choice of modems; the range of choice will be limited.

In most European nations, the PTT is interested only in the device that connects to the end of its telephone circuit; in the case of the MINET, this is always a modem. However, both the British and German PTTs in some cases also required that equipment on the digital side of the modems receive Host Nation Approval. For example, remote terminals connected to TACs over Deutsche Bundespost links will require Host Nation Approval.

4.1.2 DCS Circuits

Most of the existing DCS circuits in the MINET testbed region are already allocated. An existing circuit can be freed and made available for a new use, but only if the new use is of higher priority. Logistics applications have relatively low priority, according to DCA policy, making reallocation of existing circuits unlikely. Moreover, new circuits are not likely to be made available to the movements community in the time frame required by the MINET testbed.

Those DCS circuits that are able to be acquired for the MINET testbed should be free from preemption. Unlike voice circuits in which one call can be preempted for a higher priority one, full availability, barring malfunctions, should be guaranteed for the MINET trunk circuits.

4.2 Hardware

4.2.1 Terminals

The standard terminal to be used in the MINET is an ASCII asynchronous terminal. A MINET terminal will connect directly to (1) a TAC, (2) an electronic mail host, (3) a cryptographic unit, or (4) a low-speed modem. In each case, the electrical interface for the terminal should be CCITT recommendation V.24 (RS-232-C). All of the MINET terminals should be set up to operate with a 220 volt, 50 Hz. power source.

There are two basic types of MINET terminals: hard copy terminals and display terminals. For both of these types, there is a wide variety of manufacturers. However, the number of manufacturers that can be considered for the MINET may be severely limited owing to the homologation requirements. As mentioned above, a MINET terminal may connect to any of four devices. At least in the case of a connection to a low-speed modem, Host Nation Approval of the terminal is required. The simplest terminal acquisition policy is to consider only

homologated terminals, even if they are not required in each of the four cases. The more specialized terminals, such as portable terminals, terminals with more reliable communication protocols, and terminals with bulk media input devices, are not available from as many sources. Homologated versions of these products may be difficult to find.

Other important considerations in the choice of terminals are reliability and serviceability. To increase the likelihood of good service support, only terminals provided by local companies or foreign firms with a well-established European support staff should be considered.

As part of this study, Siemens and Perkin-Elmer have been identified as companies that sell and service terminals which have received Host Nation Approval in many western European nations. The Siemens PT-80 is a possible choice for the standard MINET hard copy terminal; one of the Perkin-Elmer CRTs could serve as the standard MINET display terminal.

4.2.2 Modems

There are basically two types of modems to be used in the MINET testbed: high-speed modems for the inter-IMP trunk circuits, and low-speed modems for the tail circuits that connect TACs or EM hosts to terminals. The high-speed modems, which will operate at a data rate of 9.6 Kbps, will be obtained from

manufacturers with a marketing, supply, and maintenance base in the testbed MINET countries. The low-speed modems, which operate in the range of 300 to 1200 bits per second, are available on lease from the PTTs in each of the testbed MINET countries.

The Deutsche Bundespost (the German PTT) may require that by 1985 all modems, both low- and high-speed, be leased from itself. At the present time, the Bundespost does not supply 9.6 Kbps modems.

The high-speed modems will always operate on dedicated circuits. Most of the low-speed modems will also be used on dedicated circuits, but some, such as those at the Nuremberg site, will be for dial-up circuits.

All modems that are not supplied by the PTTs must be on the Host Nation Approval list for the PTT in the country of interest. From time to time, new products receive approval and are added to the list. The lowest risk course is to specify modems that are already on the list, unless they are clearly inferior to a modem that is expected to receive approval soon. For the MINET, the Host Nation Approval question is related only to the high-speed modems, as the low-speed modems are PTT-furnished equipment. The high-speed modems recommended in this report have received Host Nation Approval in each of the testbed MINET countries.

The 9.6 Kbps modems that have received Host Nation Approval in western Europe all adhere to a data modulation scheme that is specified in the CCITT recommendation V.29. The V.29 recommendation does not cover another important aspect of modem operation, loop-back testing. Adherence to the CCITT recommendation for local and remote loop-back testing, recommendation V.54, was not required at the time that many of the 9.6 Kbps modems were approved. Consequently, most of the V.29 series modems use some ad hoc loopback standard (which in some cases provides more precise fault isolation than the V.54 recommendation). In anticipation of the trend towards European modems that will incorporate the V.54 recommendation, however, the MINET high-speed modems should also adhere to V.54. Picking a single loop-back standard also allows interchangeability of modems in the testbed MINET, so that a second-sourcing strategy can be followed. Modifications to existing modems to meet V.54 are straightforward enough that the resulting product is practically as low-risk as off-the-shelf equipment. The electrical interface to the digital side of the modem should be CCITT recommendation V.24 (RS-232-C).

Both Codex, a subsidiary of Motorola, and Racal-Milgo manufacture V.29 series modems. Both companies sell and support 9.6 Kbps modems in western Europe. During this study, only the Codex product has been examined carefully. It is approved in all

the MINET host nations, and a V.54 version of the Codex V.29 modem has already been developed for another customer.

4.2.3 Cryptographic Equipment

Cryptographic units, incorporating the DES encryption algorithm, will be used on the testbed MINET trunk circuits. It is also possible to use such units on the tail circuits that connect a remote terminal to a network node. The trunk circuit units can operate at 9.6 Kbps full duplex; that is, they can encrypt and decrypt 9.6 Kbps data streams in parallel. These units must incorporate knowledge of the ARPANET IMP-IMP link level protocol used on the trunk circuits. Since the link level of the IMP-IMP protocol is a variant of the BSC synchronous protocol, modifications must be made to any existing units for MINET (or ARPANET) use. These modifications, however, involve only a firmware change to the units. In the case of the tail circuits, which will employ standard low- and medium-speed link level protocols, off-the-shelf units are available.

The electrical interface on both sides of the cryptographic units will be CCITT V.24, or equivalently, RS-232-C. In order that automatic local and remote loop-back tests can take place, the V.24 interface on the cryptographic unit must pass the test signals specified in the V.54 recommendation.

It is likely that the German PTT, the Deutsche Bundespost, will require that the cryptographic units in the MINET testbed receive Host Nation Approval. In the other MINET host nations, Host Nation Approval does not appear to be necessary.

Both Codex and Racal-Milgo manufacture cryptographic units that provide link encryption for the standard low- and medium-speed link level protocols, and in particular, for the 9.6 Kbps BSC synchronous protocol. Only the Codex product has been investigated carefully in this study. The Codex unit, with a straightforward firmware modification, can be configured to work with the IMP-IMP protocol. Furthermore, the Codex unit is in the process of receiving Host Nation Approval in Germany.

4.2.4 Electronic Mail Hosts

As indicated earlier in this report, the EM hosts in the MINET testbed will support the UNIX operating system. The EM hosts will be large minicomputers, configured to operate with 220 volt, 50 Hz. power. They must be located in a computer room.

For purposes of this study, the required main memory capacity has been estimated to be 1 megabyte, and the required secondary memory capacity to be about 300 megabytes. These estimates are conservative, and lead to budgetary cost estimates that may be higher than necessary. When the implementation of the testbed MINET begins, a more precise storage estimate can be

made, on the basis of network mail traffic volume, the mail holding time, and the particular electronic mail software that is selected.

Both DEC and BBNCC manufacture systems that can serve as EM hosts. Systems that have been considered in this study are the DEC PDP 11/44, the DEC VAX 11/780, and the BBNCC C/70. The VAX 11/780 is the most expensive of these three, and the C/70 is the least expensive. At this time it is premature to recommend a particular hardware configuration; the configuration choice can be made when the MINET testbed implementation effort begins. At that time, the prices of these products and newer products such as the DEC VAX 11/750, can be reviewed and a decision made on the basis of cost, local service, and delivery availability.

4.2.5 Network Control Center Hosts

The primary network control center for the MINET testbed will be located at one of the testbed sites in Germany. (For backup purposes, the network control center for the ARPANET can also support the MINET testbed.) It is necessary to put the network control center host in a computer room. The floor space requirements, however, are small. The network control center host will run on a 220 volt, 50 Hz. power source.

The network control program runs under the UNIX operating system. One host is dedicated to running the NCC program. The

NCC program requires a dedicated large minicomputer. The amount of primary and secondary memory required depends on the size of the network to be monitored; in the case of the MINET testbed 256K bytes of primary memory and 80 megabytes of secondary storage are adequate, and represent a somewhat conservative estimate.

Both DEC and BBNCC can provide configurations for the NCC host. Likely choices for the NCC hardware include the BBNCC C/70, the DEC 11/44, and the DEC VAX 11/750. When the MINET testbed implementation is underway, the NCC host costs should be reviewed to determine the most cost-effective off-the-shelf host.

4.2.6 Interface Message Processors

The IMPs for the MINET testbed will be standard ARPANET IMPs with minor software modifications, such as V.54 loopback control, for the MINET environment. The IMPs do not need to be in a computer room, although they should be in a location with temperature and humidity ranges no worse than an office, for all 24 hours of the day. The testbed MINET IMPs will be set up to use 220 volt, 50 Hz. power. The IMP hardware will be BBNCC C/30 minicomputers.

It is likely that, at least for the German sites, the IMPs must be certified by the PTT. This certification process is simpler than the homologation process -- it is similar to

Underwriter's Laboratories approval in the United States. The Deutsche Bundespost should be contacted as soon as the implementation program is underway, to explore the certification requirements.

4.2.7 Terminal Access Controllers

The TACs for the MINET will be standard ARPANET TACs. Each TAC will be collocated with an IMP. Typically, these two units will be mounted together in a single hardware rack. The TAC will operate on a 220 volt, 50 Hz. power source. Like the IMPs, the TACs are BBNCC C/30 minicomputers.

4.2.8 Internet Gateways

Current plans call for a single internet gateway between the MINET and the ARPANET. This gateway, which will probably be located at Ft. Dietrick in the CONUS, will be a standard gateway for interconnecting two ARPANET-like networks. The gateway, which is available from BBN, is based on DEC LSI-11 hardware.

4.2.9 Hardware for Europe-CONUS Link

The 50 Kbps satellite link that will connect the MINET to the ARPANET will use modems and cryptographic equipment that are different from those used by the intra-theater MINET links. The wideband modems, as well as the cryptographic units, will be Government Furnished Equipment (GFE). Depending on the type of

GFE cryptographic equipment that will be available, it may be necessary to use a Crypto Ancillary Unit (CAU) at each end of the Europe-CONUS link. The two C/30 IMP ports on either end of the Europe-CONUS link will conform to a different electrical interface standard than V.24, which will be capable of supporting the 50 Kbps/second data rate.

4.3 Hardware Sites

All of the hardware for the MINET testbed -- hosts, IMPs and TACs, and terminals -- will be located at U.S. military installations. Most of the computer hardware can operate unattended. Nevertheless, all of the rooms which contain any of the testbed computer hardware must be accessible 24 hours a day, 7 days a week, in the event of a suspected hardware failure or some other contingency. The access to these rooms can be carefully controlled, and it is assumed that such controls will exist. However, it must be possible for an authorized individual to gain access to the computer hardware at any time.

All hardware, both computers and terminals, should be configured to operate using the 220 volt, 50 Hz. power supply. It will then be unnecessary for any of the testbed MINET hardware to depend on power converters, which tend to be unreliable.

The locations where the NCC and electronic mail hosts are installed are the most critical ones for site planning. The NCC

and EM hosts must be installed in rooms with adequate air conditioning and power supplies. Much of the heat produced by this equipment is due to the rotating mass storage devices -- the disk and tape drives.

To provide a guide to the amount of air conditioning required, power and cooling requirements for three possible EM hosts are included here. The host configurations used for the purposes of these estimates all have 1 megabyte of main memory, 2 disk drives of approximately 150 megabytes each, and 1 tape drive.

1. DEC VAX 11/780 EM host: power consumption is about 12,400 watts, yielding about 42,200 BTUs of heat generated per hour.

2. DEC PDP 11/44 EM host: power consumption is about 8000 watts, yielding about 27,200 BTUs of heat generated per hour.

3. BBNCC C/70 EM host: power consumption is about 8800 watts, yielding about 29,900 BTUs of heat generated per hour.

Sufficient air conditioning must be available to maintain a comfortable temperature for an office. Staying within the range of 65 to 75 degrees Fahrenheit would be ideal.

The other computer hardware -- the IMPs, TACs, and gateway -- may not require special air conditioning equipment, since they produce much less heat. Nonetheless, their sites must remain

within comfortable temperature and humidity ranges both day and night. The terminals are expected to be in offices, and are designed to operate in such a temperature and humidity range.

4.4 Testbed Network Operation

This section describes an organizational structure to support the operation of the MINET testbed. In addition, the scheme for troubleshooting the trunk circuits and for loading cryptographic keys is presented.

4.4.1 Responsible Organizations

Having obtained conditioned trunk circuits, it is necessary to be able to determine whether any circuit falls below the specified quality standard. Test procedures based on loop-back techniques, described later in this section, can help isolate substandard circuits. The MINET operations staff must also have line quality test equipment to measure the parameters of the suspected circuit. In this way, the PTT or DCS representative can be told specifically how the circuit is failing to meet its specification.

The organization with overall responsibility for the day-to-day operation and maintenance of the testbed MINET will be DCA Europe. DCA Europe, in order to provide this network control service for the testbed, will designate an individual or office

to be the central contact point for operations. The NCC will be at one of the testbed MINET sites. If possible, the office of the operations manager should be at the same site as the NCC.

Contacts at each of the individual PTTs serving the testbed MINET should be established to respond to reports of trouble on the circuits. The contractor for the MINET testbed will establish an office in Germany, headed by the contractor's network liaison, to which reports of hardware trouble can be directed.

Operations problems in the testbed are generally detected by the NCC, and can often be corrected without causing severe degradation of the network service. In some cases, network users may discover operations problems, in which case they should contact the operations manager. Once the operations manager and the NCC are aware of a difficulty, they can begin to isolate the fault. If possible, the NCC should determine whether the fault is in a circuit, or in some of the network hardware. In the case of suspected circuit difficulties, the NCC should contact the appropriate PTT representative. For suspected hardware trouble, the NCC should contact the contractor's support office.

4.4.2 Troubleshooting MINET Circuits and Hardware

A method for troubleshooting the MINET testbed communication hardware and circuits, especially the trunk circuits, is essential for the success of the testbed program. An effective method for troubleshooting communication links is to set up a loop-back test, in which data coming from one direction on a communication path is reflected back to its source. In the testbed MINET, a "source" IMP can send out a test data pattern, which can be looped back at any of a number of possible points between the source IMP and the "target" IMP that is at the other end of a trunk circuit. This method is used to isolate faults in the communication path.

The CCITT Recommendation V.54 specifies a loop-back scheme, which is the one recommended for use in the testbed MINET. This scheme is more likely than others to receive approval by the various nations, and makes it possible to involve modems from different manufacturers in the loop-back tests. The V.54 Recommendation specifies 4 loops: a local DTE (IMP) loop, a local DCE (modem) loop, a remote DCE (modem) analog loop, and a remote DCE (modem) digital loop.

Two potential difficulties that may hamper an effective circuit test plan are (1) the existence of cryptographic equipment in the network, and (2) the policies of the various

PTTs in the MINET host nations. The first potential problem, which is a technical one, can arise if the cryptographic equipment does not pass loop-back control signals from the IMP on to the modem, or if there is not a convenient way to re-establish cryptographic synchrony when running a loop-back test. These technical problems do not arise in the case of the Motorola Infoguard unit, which does pass through the loop-back signals, and which has a self-synchronizing cipher. The second problem, which is an administrative one, is that some of the PTTs -- the Deutsche Bundespost in particular -- may not allow certain kinds of loop back tests, or may not allow some kinds of test equipment to be attached to their circuits. For example, the Bundespost may not allow the V.54 remote DCE analog loop. However, the local loops together with the remote DCE digital loop still provide a good fault isolation capability.

It will be possible to control all of the loop tests from a terminal connected to the NCC host. The NCC host can instruct a selected IMP to run various loop tests. Thus no on-site personnel are required to do loop-test troubleshooting of communication paths. If the communication path between two IMPs "A" and "B" fails, the NCC operator can run loop-back tests from IMP "A," trying successively larger loops until IMP "A" fails to receive its own test pattern. The NCC operator can then run a similar sequence of tests from IMP "B." The combination of the

information returned from the IMP "A" and IMP "B" loop-back tests can in most cases isolate the fault to an IMP and/or cryptographic unit, a modem, or the trunk circuit. Having located the faulty component, or at least limited the extent of the fault, the NCC operator can contact DCA Europe if a circuit problem is suspected, or the European office of the contractor if a hardware problem is suspected.

4.4.3 Cryptographic Key Loading

The DES-based cryptographic units, such as the Motorola Infoguard unit, generally use a two-level keying hierarchy with primary and secondary keys. The secondary keys, which protect the actual data on the trunks, can be down-line loaded and encrypted under the primary keys. In this down-line loading process, one DES unit acts as a master and the other as a slave. An operator loads the slave unit from the master unit. Re-loading of secondary keys should be planned to take place when an interruption in service, on the order of a minute or two on a per line basis, can be tolerated. The network will automatically reroute traffic around such single trunk outages. The new secondary keys loaded by the operators can be generated and initialized independently. The rate and timing of secondary key loading will be controlled by administrative policy. This loading does not have to be synchronized by line or by site. As a matter of good practice, the operators should load different

secondary keys into each of their master units. Of course, each slave unit must be loaded with the same secondary key as its corresponding master unit. As indicated in subsection 3.5, master cryptographic units at London, Bremerhaven, Frankfurt, Stuttgart, Sigonella, and Athens will eliminate the need for physical transport of secondary keys between any of the testbed MINET sites. The network management office will need to issue proper operation, logging, and storage guidelines for the key generation and loading devices in order to safeguard data protection.

4.5 Development and Testing Activities

This section describes the development and testing activity, for both hardware and software, that will take place prior to the installation of the testbed MINET. All of the development activity and much of the testing activity will take place at the MINET testbed contractor's site in the CONUS. Given the objectives of the MINET testbed, all of these development activities are low-risk activities.

4.5.1 Hardware Development and Testing

The cryptographic unit for the trunk circuits will require modification in order to accommodate the IMP-IMP link level protocol. The off-the-shelf DES units considered in this study encrypt at the link level, and so must take the link level

protocol into account. No commercial products are designed to work with the IMP-IMP link level protocol. However, some of the products, including the Motorola Infoguard, are designed to work with the BSC link level protocol, which is similar to the IMP-IMP link level protocol. The Infoguard unit can be modified for the IMP-IMP protocol for less than $10,000.

Depending on the modem that is chosen for the trunk circuits, it may be necessary to make a small hardware modification so that the modem can receive the V.54 loop-back signals at its electrical interface. Newer modems are more likely to incorporate this automatic V.54 loop back feature.

Following the acquisition of the modified DES units and the modems with automatically controlled V.54 loopback, a test network link, consisting of an IMP, DES unit, and modem on each end connected by a telephone circuit, will be set up in the contractor's staging area. Using this test link, the various loop-back tests will be performed to ensure that all components work as expected. Once all of the components have passed these tests, the remaining components produced for the testbed can be shipped directly to the theater.

4.5.2 Software Development and Testing

The IMP software will require a number of modifications for the MINET environment. In order to run the loop-back tests, the IMP software must be modified to take into account the existence of the DES cryptographic units. Modifications must also be made to allow the IMP to activate each of the test loops defined in the CCITT V.54 recommendation. The software that determines whether a line is up or down must be modified to deal with the 9.6 Kbps trunk circuits.

The software for the NCC host must be also be modified in order to direct the IMPs to run the V.54 loop-back tests. Current NCC host software can control loop-back tests, but these tests are less extensive than the V.54 loop-back tests.

The EM software development required depends on whether MH or Infomail is chosen as the the MINET mail facility. If Infomail is chosen, little tailoring of the user interface will be required. If MH is chosen, major changes to the user interface will be required to create an interface more acceptable to the inexperienced computer user. In either case, forms needed by the testbed MINET users must be incorporated into the mail system. This study makes a worst-case assumption for the level of EM software manpower required, which is that MH is chosen and that substantial user interface modifications as well as some more basic structural modifications will be necessary.

4.6 Manpower Plan for Operational Support

The staff supplied by the contractor for operational support in the theater will be headed by the contractor's network liaison. The network liaison's staff will consist of a programmer, a training support representative, two hardware maintenance representatives, and operators for the EM and NCC hosts. The liaison will arrive in the theater in the third month after the testbed implementation begins, and the rest of the staff will arrive beginning in the sixth month.

The network liaison will be responsible for establishing an overseas office for the contractor, probably near the NCC site. The liaison will work closely with the DCA Europe network operations manager in setting up an NCC site. The contractor will be responsible for providing the NCC operators. As a cost saving measure, local nationals will be hired as NCC operators wherever possible. Each EM site is assumed to have one operator.

To reduce manpower costs in years 3 and 4 of the testbed operation, USEUCOM may choose to phase out the operations staff supplied by the contractor, replacing it with military personnel. It may also be possible to replace other members of the contractor's overseas staff with military personnel, depending on the availability of q alified individuals.

The testbed MINET is planned to operate on a continuous basis, with occasional scheduled interruptions in network service for hardware preventative maintenance, software updates, and cryptographic key loading. Operator support, however, will be available at the EM sites only at the following times:

NCC/EM site -- operator support 0800 to 2400, Monday through Friday, with two operators on duty during this period;

Other EM sites -- operator support 0800 to 1700, Monday through Friday, with one operator on duty during this period.

The main role of the [illegible] host operators will be to perform file backup and recovery operations. For typical composition and reception of electronic mail, operator support at the EM host is not required.

4.7 Budgetary Cost Estimates

This section presents the budgetary cost estimates for the hardware, circuits, manpower, and miscellaneous components of the MINET testbed. For each of these four components, costs are broken down according to the four years of the testbed MINET program. The cost estimates shown here are essentially those presented to the General Officer's Steering Committee in August 1980. Since then, USEUCOM has further developed these cost estimates for the purpose of DoD planning and funding requests.

4.7.1 Costing Assumptions

There are some costs which have not been included in the budgetary estimate of this study. The important assumptions underlying the budgetary cost estimates are listed here.

1. The cost of providing access to the MINET for the data base hosts is not included in this budget. The assumption made in the case of the testbed MINET, which is consistent with the DCA policy governing the ARPANET, is that interfacing costs are one of the costs of operating the host systems. Based upon the experiences of various ARPANET hosts, a rough estimate is that the interfacing costs for the four data base hosts will be, at most, 10% of the total 14.5 million dollar budgeted MINET testbed cost, or about 1.5 million dollars. It may be possible, however, to attach all four hosts for a much lower cost. A more precise estimate of these costs involves a more thorough study of the state of existing applicable software and hardware in the ARPANET community.

2. This budget plan assumes that all hardware will be purchased, rather than leased. All circuits, however, will be leased.

3. The contractor will bid a turnkey cost plus fixed fee system contract, with level of effort support in the theater.

4. The contractor-supplied documentation will be of good commercial quality but will not comply with Mil-Std practices.

5. Inflation is not taken into account; these estimates are in terms of constant 1980 dollars.

6. The cost of any terminals in the CONUS that will connect to the testbed MINET (through the ARPANET) has not been included.

7. The costs of maintaining a field office for the contractor have not been included. It is assumed that some office space at one of the central Germany testbed MINET sites can be made available to contractor personnel.

8. The costs of a staging facility in the CONUS, where the hardware and software for the testbed will receive its initial testing and debugging, has not been included. An example of a staging cost is the cost of computer time needed for the testing.

9. The costs involved in preparing the EM host, IMP/TAC, or terminal sites for the testbed have not been included, except for the cost of setting up the NCC site.

4.7.2 Hardware Costs

The budgetary cost estimates for the MINET testbed hardware are given in Table 4-1. A brief description of each item listed in Table 4-1 is included here.

1. EM host -- three electronic mail hosts are budgeted for the MINET testbed program, with two acquired in stage 1 and one acquired in stage two. For this budgetary estimate, the least expensive of the C/70, VAX 11/780, and PDP 11/44 systems is assumed. This system is the BBNCC C/70, which consists of 1 CPU with 1 megabyte of memory, one operator terminal, 2 160-megabyte disks, 1 8-line EIA terminal interface, one 1822 interface for network access, and one tape drive. The estimated price is $104,000.

2. NCC host -- one NCC host is budgeted for acquisition in stage 1. The hardware used for the budgetary estimate is a BBNCC C/70 system, comprising 1 CPU with 256K bytes of memory, one operator terminal, 2 80-megabyte disks, 1 8-line EIA terminal interface, 1 1822 interface for network access, and 1 tape drive. The estimated price is $89,700.

3. IMP -- the IMP hardware is a BBNCC C/30 computer, which will cost $21,500.

4. TAC -- the TAC hardware is a BBNCC C/30 computer, costing $23,000.

5. Gateway -- the internet gateway host is a DEC LSI-11 system with 1822 interfaces for network access. The cost is $19,800.

6. DES units -- the DES unit used for the budgetary cost estimate is the Motorola Infoguard unit. Normally these units will be packaged two per rack at the various MINET testbed nodes. There will also be a modest extra engineering charge from Motorola to modify the Infoguard firmware to accommodate the IMP-IMP protocol. The estimated cost for a modified Infoguard unit, plus one half of a rack, is $4000.

7. DES Key Set -- the DES Key Set is used to generate cryptographic keys and load them into the DES units. For budgetary estimates, the key set from the Motorola Infoguard line is used. One key generator plus one key loader will cost approximately $7200.

8. CAU -- two Crypto Ancillary Units may be needed, depending on the type of GFE cryptographic equipment used on the Europe-CONUS link. For the budgetary estimate, units from Dataproducts, New England are assumed. The price is $4400 each.

9. High-Speed Synchronous Modem -- these modems will be used on all the trunk circuits except the Europe-CONUS trunk. For the budget estimate, Codex LSI 96/V.29 modems are assumed. A small surcharge may be added by Codex to cover the cost of providing a standard V.54 loopback interface. The estimated cost is $7200 per modem.

10. Asynchronous Modem -- these modems will be used to connect remote asynchronous terminals to TACs. These lower speed modems generally must be acquired from the PTTs. In a number of instances, explained in the section on circuit costs, it may be possible to multiplex circuits from TACs to terminals in order to save on circuit lease costs. Of the 24 stage 1 terminals that are planned to connect to TACs over leased tail circuits, only 11 are assumed to have dedicated circuits; the other 13 terminals will share multiplexed circuits. For the multiplexed circuits, medium-speed synchronous modems will be used. There are thus 22 asynchronous modems planned for stage 1. The estimated price of a low-speed modem is $700.

11. Dial-in Modem -- in order to support dial-up terminal access to TACs, dial-in modems must be acquired, and generally are available only from the PTTs. The estimated cost is $900 each.

12. Medium Speed Synchronous Modem -- for those areas in which remote terminals must connect to TACs over unreliable voice-grade circuits, synchronous modems which can operate over non-conditioned voice-grade circuits can be used. There are nine sites, all of which are scheduled to join the testbed in stage 3, which will connect to their respective TACs over tails that are presumed to be unreliable. These sites are Lisbon, Torrejon, La Maddalena, Coltano, Souda Bay, Izmir, Iskenderun, Incirlik, and Cakmakli. In addition, for those sites where circuit costs can be reduced by multiplexing lines from terminals to TACs, a medium-speed modem is required. There are five pairs of stage 1 sites where multiplexed tail circuits are assumed: Kaiserslautern to Ramstein, Mannheim to Heidelberg, Rhein/Main to Frankfurt, Mildenhall to London, and Nuremberg to Heidelberg. The Codex MX 2400 modem is used for the budgetary estimate, with a price of $1700.

13. Asynchronous-Synchronous Statistical Multiplexer -- these devices, which can support a small number of asynchronous terminals while using a synchronous protocol on the network side, are planned to be connected to the medium-speed synchronous modems. They can serve two distinct purposes -- to support terminals that connect over unreliable circuits, and to multiplex circuits to save line costs. The Codex 664

Statistical Multiplexer, used for budgetary purposes here, incorporates an error detection and retransmission scheme to achieve reliable transmission. It can support up to 4 asynchronous terminals, any of which can operate at 300, 1200, or 1800 bits per second. Ten units of this type are planned in stage 1 for the multiplexed tail circuits, and 18 are planned in stage 3 for the unreliable tail circuits. The budgetary cost estimate is $2300.

14. Host port -- an IMP has two kinds of ports: host ports, which connects to hosts, and trunk ports, which connect to the modems. For each testbed MINET IMP except the one at the CONUS end of the Europe-CONUS link, a host port for the nearby TAC is required. In addition, host ports are needed for the DAMMS host, for the EM hosts, and for the NCC host. The cost is $350.

15. Trunk port -- each IMP requires a separate trunk port for each attached network trunk. The cost is $350.

16. Terminal port -- a TAC requires one terminal port for each group of up to eight attached terminals. The cost is $800.

17. Hard Copy Terminal -- for budgetary cost estimate purposes, a hard copy terminal with an impact printing mechanism is assumed. All 9 of the stage 3 terminals that are planned to connect to TACs over unreliable circuits are assumed to be hard copy terminals. The estimated price is $3000.

18. Display Terminal -- a simple display terminal with no editing features has been assumed for the budgetary cost estimate. The price is $1000.

19. Portable Terminal -- a portable, hard copy terminal, with local editing capabilities and a built-in acoustic coupler are assumed for the cost estimate. The budgeted price is $4000.

20. Bulk Input Terminal -- bulk input terminals, which are designed to accept input such as punched cards or magnetic tapes, will be needed at the air and water terminals in the theater as sources of input to the DAMMS data base. The estimated price is $12,000.

21. Racks & Cabinets -- racks and cabinets will be required for each MINET testbed node, for each EM host, and for the NCC host. Estimated cost per installation is $2000.

22. Cables -- cables will be required to connect the hardware mentioned above. The estimated cable count is two per terminal, plus one for each node and one for each host. The price is $100.

23. Test Equipment -- test equipment must be acquired for troubleshooting network hardware and circuitry. One set of equipment is assumed for stages 1 and 2, and a second set for

stage 3. The estimated cost is $50,000 for each set of test equipment.

The hardware costs, totaling 2.3 million dollars, are spread out over the MINET testbed program as follows: all stage 1 costs are incurred in year 1, all stage 2 costs in year 2, and all stage 3 costs in year 3. No hardware costs are budgeted for year 4 of the testbed program.

4.7.3 Circuit Costs

This section provides the cost estimates for the testbed trunk and tail circuits. Trunk circuits are those that connect IMPS, and tail circuits are those that connect a remote terminal to a TAC. Circuits for the MINET testbed will be obtained either from the PTTs or from DCA. DCA circuits can be provided to the testbed MINET at no cost, while PTT circuits must be leased. The cost estimate made in this section assumes a network with all PTT circuits except for the Germany-CONUS and the Turkey-Germany DCA circuits. It thus represents an upper bound on the testbed MINET circuit costs, since it is likely that the testbed will contain some additional DCA circuits in the theater.

Since the testbed spans a number of nations, the pricing schedules for the domestic circuits in each nation, as well as pricing schedules for the international circuits, are needed for a precise cost estimate. The method used in this report is an

Table 4-1

Hardware Budgetary Estimate for the MINET Testbed
Prices in Thousands of dollars

			STAGE 1		STAGE 2		STAGE 3	
		K$/unit	qty	cost	qty	cost	qty	cost
1.	EM Host	104.0	2	208.0	1	104.0	-	---
2.	NCC Host	89.7	1	89.7	-	---	-	---
3.	IMP	21.5	8	172.0	3	64.5	2	43.0
4.	TAC	23.0	7	161.0	3	69.0	2	46.0
5.	Gateway	19.8	1	19.8	-	---	-	---
6.	DES Units	4.0	16	64.0	8	32.0	6	24.0
7.	DES Key Set	7.2	4	28.8	1	7.2	1	7.2
8.	CAU	4.4	2	8.8	-	---	-	---
9.	High-Spd Synch Modem	7.2	16	115.2	8	57.6	6	43.2
10.	Asynchronous Modem	0.7	22	15.4	12	8.4	-	---
11.	Dial-in Modem	0.9	8	7.2	-	---	-	---
12.	Med-Spd Synch Modem	1.7	10	17.0	-	---	18	30.6
13.	Statistical Mux	2.3	10	23.0	-	---	18	41.4
14.	Host Port	0.35	11	3.9	4	1.4	2	0.7
15.	Trunk Port	0.35	18	6.3	8	2.8	6	2.1
16.	Terminal Port	0.8	12	9.6	6	4.8	2	1.6
17.	Hard Copy Terminal	3.0	44	132.0	25	75.0	14	42.0
18.	Display Terminal	1.0	10	10.0	4	4.0	2	2.0
19.	Portable Terminal	4.0	10	40.0	-	---	-	---
20.	Bulk Input Terminal	12.0	5	60.0	-	---	-	---
21.	Racks & Cabinets	2.0	10	20.0	4	8.0	2	4.0
22.	Cables	0.1	135	13.5	74	7.4	14	1.4
23.	Test Equipment	50.0	1	50.0	-	---	1	50.0
	Subtotal			1275.2		446.1		339.2
	add 10% for spares			127.5		44.6		33.9
	Stage Totals			1402.7		490.7		373.1

approximation, in which all circuits, domestic or international, are priced using the circuit cost algorithm of the German PTT, the Deutsche Bundespost. Based upon the experience of DCA Europe, the Bundespost lease rates are generally as high or higher than the domestic rates of other PTTs as well as the international rates. Thus, using the Bundespost rates in every case will produce an upper bound on the estimated circuit lease costs. The Deutsche Bundespost pricing algorithm for circuits is:

- Basic 2-wire circuit -- 20 Marks/Kilometer/Month
- Surcharge for a 4-wire circuit -- 400 Marks/Month
- Surcharge for M1020 conditioning -- 480 Marks/Month
- Installation charge for a 4-wire circuit -- 1600 Marks

It is assumed that all trunk and tail circuits are 4-wire, conditioned circuits. This assumption is also conservative, since tail circuits may not need to be 4-wire or conditioned circuits. Based on this assumption, the monthly charge for any testbed MINET circuit will be 880 Marks plus 20 Marks per kilometer and the installation charge will be 1600 Marks. All circuit prices will be given in dollars, using the conversion rate of 1 Mark = 0.58 dollars.

The budgetary cost estimates for the MINET testbed circuits are given in Table 4-2. In stage 1, only 16 tail circuits are

Table 4-2

Circuit Budgetary Estimate for the MINET Testbed
Prices in Thousands of Dollars Per Month

		Circuit	KM.	unit cost	qty	total cost
Stage 1	Trunks	LON-RDM	300	4.0	1	4.0
		RDM-BRM	350	4.6	1	4.6
		BRM-FRK	390	5.0	1	5.0
		FRK-HDL	80	1.4	1	1.4
		HDL-STT	80	1.4	1	1.4
		STT-RAM	130	2.0	1	2.0
		RAM-FRK	70	1.3	1	1.3
		RAM-LON	650	8.1	1	8.1
		RAM-CONUS	(DCS)	---	1	---
	Tails	HDL-Nuremberg	180	2.6	1	2.6
		short tails	---	1.0	15	15.0
	Total circuits and cost				25	45.4
Stage 2	Trunks	LON-ROT	1700	20.0	1	20.0
		ROT-SIG	1950	23.1	1	23.1
		SIG-NAP	450	5.7	1	5.7
		NAP-STT	1000	12.1	1	12.1
	Tails	NAP-Gaeta	80	1.4	1	1.4
		short tails	---	1.0	1	1.0
	Total circuits and cost				6	63.3
Stage 3	Trunks	NAP-ATH	900	11.0	1	11.0
		ATH-IST	600	7.5	1	7.5
		IST-Germany	(DCS)	---	1	---
	Tails	LON-Edzell	600	7.5	1	7.5
		LON-Holy Loch	550	6.9	1	6.9
		ROT-Lisbon	350	4.6	1	4.6
		ROT-Torrejon	475	6.0	1	6.0
		NAP-La Maddalena	425	5.4	1	5.4
		NAP-Coltano	475	6.0	1	6.0
		ATH-Souda Bay	275	3.7	1	3.7
		IST-Izmir	350	4.6	1	4.6
		IST-Iskenderun	800	9.8	1	9.8
		IST-Incirlik	700	8.6	1	8.6
		IST-Cakmakli	250	3.4	1	3.4
	Total circuits and cost				14	85.0

assumed, even though there are 24 terminals that could connect to TACs or EM hosts by short, low-speed tail circuits. This savings of 8 tail circuits is accomplished by assuming some multiplexing:

1. 5 Kaiserslautern terminals sharing 3 circuits to Ramstein, with 2 dedicated circuits and 1 3-way multiplexed circuit;

2. 2 Mannheim terminals sharing 1 circuit to Heidelberg, with 1 2-way multiplexed circuit;

3. 3 Rhein/Main terminals sharing 2 circuits to Frankfurt, with 1 dedicated circuit and 1 2-way multiplexed circuit;

4. 6 Mildenhall terminals sharing 3 circuits to London, with 2 dedicated circuits and 1 4-way multiplexed circuit; and

5. 2 Nuremberg terminals sharing 1 circuit to Heidelberg, with 1 2-way multiplexed circuit.

Each dedicated circuit that is saved eliminates 2 low-speed asynchronous modems. On the other hand, each shared circuit requires 2 medium-speed modems (listed as hardware item 12) and 2 statistical multiplexers (listed as hardware item 13). The count of the items in the hardware list incorporates these multiplexing assumptions. The net extra hardware cost to provide the multiplexing is about $22,000, while the circuit cost savings is about $9600 per month. Thus there is a payback period of less than 3 months. These multiplexing assumptions are conservative,

since it is possible to reduce even further the number of tail circuits by relying to an even greater extent on multiplexing equipment.

Table 4-2 shows the monthly costs for each of the three testbed MINET stages. In addition to the monthly circuit costs, there are the one-time charges for circuit installation. The charge for a 4-wire circuit is 1600 Marks, or 928 dollars, assuming 1 Mark = 0.58 dollar. Thus the installation charge for 24 stage 1 circuits (no charge is assumed for the DCA circuit) is $22,300; the stage 2 installation is $5600; and the stage 3 installation is $12,100.

Finally, these circuit costs must be related to the four year schedule of the MINET testbed program. It is assumed that the stage 1 circuits will be obtained in the tenth month of year one, the stage 2 circuits will be obtained in the tenth month of year two, and the stage 3 circuits will be obtained in the tenth month of year three. Thus, the estimated circuit costs for the four program years, in thousands of dollars, are:

Year 1:	22.3 + 3 X 45.4	=	158.5
Year 2:	5.6 + 12 X 45.4 + 3 X 63.3	=	740.3
Year 3:	12.1 + 12 X 45.4 + 12 X 63.3 + 3 X 85.0	=	1571.5
Year 4:	12 X 45.4 + 12 X 63.3 + 12 X 85.0	=	2234.4
Total:		=	4794.7

The total estimated circuit cost for the whole program is 4.8 million dollars. The 4.8 million dollar estimate is based on the assumption that only two circuits -- the satellite circuit from Germany to CONUS and the satellite circuit from Germany to Turkey -- will be provided by DCA, and that the rest must be obtained from the PTTs. It further assumes that all PTT circuits are priced according to the Deutsche Bundespost circuit price structure, using the conversion rate of 1 German Mark = 0.58 in U.S. dollars.

4.7.4 Manpower Costs

The manpower costs for the four years of testbed operation are derived in this section. The cost estimates are based on the assumption that the contractor will be responsible for development, installation, operation, and much of the maintenance of the MINET testbed. These cost estimates can be reduced substantially if military personnel assume operational responsibility for the MINET testbed in the third and fourth years of the program. The estimates presented here, based on contractor operation, thus serve as an upper bound on the cost of contractor-supplied manpower.

The manpower assignments are calculated in terms of the program years, rather than in terms of the three program stages. The first year requires a modest development commitment in the

CONUS and a corresponding build-up of staff in Europe. The second through fourth years require only minimal CONUS support, but an 11-member European staff. The budgetary cost estimates for the MINET testbed manpower requirements are given in Table 4-3. The CONUS staff for the testbed implementation will include an individual serving as the manager and overall design engineer, programmers for the electronic mail enhancements, programmers for the IMP and NCC software extensions, and a documentation and training representative. The CONUS staff can be reduced to just the manager after the first program year. The European staff will consist of a network liaison, a programmer, a training support representative, two hardware maintenance representatives, and operators for the EM and NCC hosts. In the case of the CONUS staff, it will be possible to use "fractions" of people; the IMP and NCC programmers are budgeted on a half-time basis. The European staff will be devoted fully to the MINET testbed; fractional person-years in this case simply indicate that the staff member begins his assignment part way through the year. The total estimated manpower costs for the program are approximately 6.6 million dollars.

4.7.5 Miscellaneous Costs

Miscellaneous costs are itemized in this section, in Table 4-4. Terminal maintenance is assumed to be a monthly charge of 2% of the purchase price of the terminals. Software license fees

Table 4-3

Manpower Requirements Budgetary Estimate for the MINET Testbed
Manpower Time Shown in Units of Person-years
Prices in Thousands of Dollars

CONUS Staff	Year	1	2	3	4
Manager / Design Engineer		1.0	1.0	1.0	1.0
EM enhancement programmers		2.0			
Hardware integration engineer		1.0			
IMP programmer		0.5			
NCC software customizing		0.5			
Documentation and training Rep		1.0			

On-site Staff	Year	1	2	3	4
Network Liaison		0.75	1.0	1.0	1.0
Programmer		0.25	1.0	1.0	1.0
Training Support Rep		0.5	1.0	1.0	1.0
Maintenance Support Rep			2.0	2.0	2.0
Operators (U.S)		0.5	2.0	3.0	3.0
Operators (Foreign)		1.5	3.0	3.0	3.0

Job Classification	Year	1	2	3	4
CONUS Professional		6.0	1.0	1.0	1.0
On-site Professional		1.5	5.0	5.0	5.0
On-site Operator (U.S.)		0.5	2.0	3.0	3.0
On-site Operator (Foreign)		1.5	3.0	3.0	3.0

Cost Estimates	Rate	Year	1	2	3	4	Total
CONUS Professional	$130		780	130	130	130	1170
On-site Professional	$175		264	875	875	875	2889
On-site Operator (U.S.)	$150		75	300	450	450	1275
On-site Operator (Foreign)	$120		180	360	360	360	1260
Total			1299	166	1815	1815	6594

will be required for copies of the UNIX operating system on the EM and NCC hosts. If the Infomail electronic mail software is used, there will be a license fee. Other EM software under consideration would not require a license. Table 4-4 includes UNIX fees and estimated Infomail license fees, all of which would fall into years one and two. Travel is broken down into overseas trips for the CONUS staff and European staffs, and travel by the European maintenance support representative to each of the testbed MINET nodes, at the rate of eight trips per node per year. The total miscellaneous costs budget for the testbed program is estimated at $800,000.

4.7.6 Total Costs

The total estimated project costs are given in Table 4-5. In this table, costs are given to the nearest $100,000. It should be emphasized that the cost estimate is in constant 1980 dollars, so inflation is not taken into account. The 14.5 million dollar MINET testbed budget estimate shown in this report was first presented at the General Officer Steering Committee meeting on August 19, 1980, at Patch Barracks. The derivation of the 14.5 million dollar estimate which is presented in this report is essentially the same as the derivation given at the GOSC meeting.

Table 4-4

Estimated Miscellaneous Costs for the MINET Testbed
Prices in dollars

Terminal Maintenance (2% of purchase price / month)

Stage	1	$4800 / month
	2	$1600 / month
	3	$2500 / month

Software Licenses

Year	1	$51000
	2	$18000
	3	-----
	4	-----

Custom Hardware Modifications

Year 1 $10000

Travel (overseas trips @ $1500, local maintenance trips @ $500)

Year	1	28 overseas trips	= $42,000
	2	8 overseas trips, 56 local trips	= $40,000
	3	8 overseas trips, 80 local trips	= $52,000
	4	8 overseas trips, 96 local trips	= $60,000

Leased Car

$3600 per year for all four years

Expendables

$10,000 per year for all four years

Site Preparation for the NCC Site and other Stage 1 Electrical Work

$200,000

Total Estimated Miscellaneous Costs

Year	1	$331,000
	2	$134,000
	3	$150,000
	4	$180,000

Table 4-5

Total Estimated Costs for the MINET Testbed
Prices in Millions of Dollars

	Program Years 1	2	3	4	Total
Hardware	1.4	0.5	0.4	---	2.3
Circuits	0.2	0.7	1.6	2.3	4.8
Manpower	1.3	1.7	1.8	1.8	6.6
Miscellaneous	0.3	0.1	0.2	0.2	0.8
Total	3.2	3.0	4.0	4.3	14.5

The final project budget estimates, for the purpose of DoD planning and funding requests, are being developed by USEUCOM. Since the GOSC meeting, USEUCOM has further developed the budget estimate given in this report, focusing on three areas. First, USEUCOM has assumed that 90% of the circuits will be PTT circuits and 10% will be DCS circuits. The estimate presented in this report is based on 100% usage of PTT circuits, except for two DCA satellite circuits. Thus, this 90%-10% assumption will reduce the circuit portion of the estimated program cost. Second, USEUCOM has increased the budgetary estimate to include funding for a System Engineering and Technical Advisory (SETA) contractor. Third, USEUCOM has increased the budgetary estimate by adding a 5% contingency cost to the hardware budget presented

here. These changes yield a revised overall estimated program cost of 16.0 million (constant) dollars.

5. RISK ANALYSIS

The testbed MINET system design and program plan presented in this report is the lowest risk cost effective approach currently possible. Many potential risks have been eliminated through careful application of packet switching and data communication technology. The program is not, however, totally free of risk. The following subsections describe the remaining risks currently identified.

5.1 Possible Changes in Traffic Projections

The network is designed based upon traffic volumes and patterns which are very unlikely to develop in reality specifically as estimated. The final Stage 3 traffic volume will not be realized until approximately 4 years after the initial traffic surveys. The basic survey data is certain to be out of date. Furthermore, the various database host interfacing options may alter the character of traffic flows. A fundamental shift in traffic flow could occur if the EM hosts are eventually not used as intermediary devices for data base inquiry. The traffic from EM hosts would shift to database hosts which are primarily in CONUS. This condition is, however, quite unlikely over the testbed period, since it would behcove the users to continue using EM service as the enquiry mechanism for reasons of responsiveness and common access interfacing.

All of these traffic changes can be accommodated by the MINET. The MINET is sized to carry a fourfold traffic increase over the estimated base. This will provide sufficient excess capacity to permit traffic change, time to notice the change, and time to react to the change. The network can easily be reconfigured to handle these traffic shifts. The system contractor will monitor system performance and suggest configuration changes as appropriate. There are no fundamental technical reasons which would prevent necessary reconfiguration other than overall limits to growth.

5.2 Limit to MINET Growth

The ARPANET has grown since its inception to a 65 node, 200 host, 10,000 terminal network. During this period service was rarely interrupted and there was no need to completely replace the fundamental hardware, software and circuit investment. The network was often expanded, reconfigured, and modified; but not replaced. The growth potential for MINET reflects the demonstrated growth of the ARPANET.

The single critical resource in the MINET which will limit such growth is the unreliable 9.6 Kbps trunk. The ARPANET enjoys a significantly better fundamental communication resource -- very reliable 50 Kbps trunks. At some point the MINET will experience severe bottlenecks due to particular overloaded trunks which will

not be relievable through reconfiguration. This will be a performance threshold which will be difficult to overcome using current ARPANET protocols. The fourfold Stage 3 traffic scenario is not far from this threshold.

A solution to this problem is to add parallel 9.6 Kbps circuits to augment the node-node capacity in those specific trunks where the problem appears. This requires the ability to treat two physical circuits as one logical circuit, while maintaining separate line up-down and fault isolation for the physical circuits.

Another solution is to modify the routing algorithm in the IMPs to permit multiple simultaneous paths between node pairs. This would permit using currently untapped bandwidth in underutilized trunks to augment the bandwidth on overutilized trunks.

Multiple circuit trunking and multi-path routing are both currently being designed at BBN specifically for the ARPANET. It is very likely that both solutions will be available on the ARPANET when the MINET will need them. These new protocols and algorithms could then be added to the MINET with modest tuning and automatic software loading into IMPs.

5.3 PTT Approval Issues

Some of the hardware selected for the MINET will require PTT Host Nation Approval. Most hardware recommended here already has such approvals in most MINET countries. Any manufacturer changes to bring the equipment in-line with PTT approvals may cause a minor change in price or a change in vendor. There exists a minor risk that the PTTs will not approve a U.S.-made packet switch to be installed on their circuits. There is also a modest risk that the PTTs will not permit the automatic interconnection of the EM service with their Telex service.

These risks can be minimized if the approvals are sought through U.S. military channels. The host nations will be much more responsive to this network and the particular interfaces if the requests are submitted to the PTT's military liaison offices. It will also have to be made perfectly clear that this is a network solely for U.S. military use. The negotiation of interface and hardware approvals or waivers will be a delicate process which should be embarked upon with extreme caution and precision so that the MINET budget or schedule is not seriously jeopardized.

5.4 Assumptions About the Capabilities of the Contractor

This design study recommends that the MINET be built by taking advantage of existing ARPANET technology. The project plan and budgetary cost estimates included in this study are based on the assumption that the MINET testbed will depend heavily on off-the-shelf ARPANET hardware and software. Only a small amount of additional development effort will be required. It is important to stress that this technology can be transferred into the operational arena of the MINET community most efficiently by a contractor with ARPANET experience.

5.5 Future MINET Requirements

This is a testbed network design study. There are issues which will surface if the network becomes operational. The long-term requirement for the post testbed network is combat readiness. This study was permitted to address the combat readiness issues only briefly. We have, however, identified several network requirements which stem from a combat readiness requirement:

- end-to-end encryption
- mobile nodes
- automatic reconfiguration of trunks
- more reliable protocols
- fully distributed mail service.

Fortunately, all of these network requirements are currently being addressed by ARPA-sponsored research. A variety of currently funded developments will be directly applicable to the operational MINET. USEUCOM and MINET management should study the combat readiness issues in order to identify specific research and development requirements. USEUCOM can then become a proponent of the appropriate research which would foster the goal of an operational MINET.

Most of the combat readiness enhancements can be added without discarding the testbed investment. The operational MINET can, therefore, evolve from the testbed MINET without the need to lose this investment in hardware, software and operational experience.

6 APPENDIX A - Tested MINET Terminal Sites

Testbed MINET Terminal Requirements

KEY: Tot = Total terminals per site; HCp = Hard Copy Terminals;
Dis = Display; Por = Portable; Rel = Reliable; Blk = Bulk Input

Location	Tot	HCp	Dis	Por	Rel	Blk
London (NAVEUR Bldg.)						
Stage 1 Sites						
NAVEUR Bldg, CINCUSNAVEUR	3	1	1	1		
Mildenhall, USEUCOM Log Cond Cell	1	1				
Mildenhall, UK Planning Cell, 4th Tr	1	1				
Mildenhall, ACA, USAFE	1	1				
Mildenhall, NAV DET CARGO	2	1				1
Mildenhall, NAV OPNS	1	1				
Felixstowe, MTMC TTGE	1	1				
Stage 1 Subtotal	10	7	1	1		1
Stage 3 Sites						
Edzell, Scotland, NAVFAC	1	1				
Holy Loch, Scotland, NAVDET	4	2	2			
Stage 3 Subtotal	5	3	2			
Total	15	10	3	1		1
Rotterdam (MTMC HQ Bldg.)						
Stage 1 Sites						
MTMC HQ Bldg, MTMC TTGE HQ	2	1	1			
MTMC HQ Bldg, MTMC BENELUX Terminal	5	1	1	2		1
MTMC HQ Bldg, TMO, 4th Trans	1	1				
MTMC HQ Bldg, USAFE WPLO	1	1				
Total	9	4	2	2		1

Location	Tot	HCp	Dis	Por	Rel	Blk
Bremerhaven (Karl Schurz Kaserne)						
Stage 1 Sites						
K. S. Kaserne, HQ 1st Mov. Region	1	1				
K. S. Kaserne, MSCEUR	6	2	4			
K. S. Kaserne, MTMC TTGE	4	1	1	1		1
K. S. Kaserne, TMO, 4th Trans	1	1				
K. S. Kaserne, USAFE WPLO	1	1				
Total	13	6	5	1		1
Frankfurt (Camp King)						
Stage 1 Sites						
Camp King, HQ 4th Trans.	4	1	1	2		
Camp King, MCC 4th Trans.	3	3				
Gibbs Kaserne, HQ 3rd Mov.	1	1				
Gibbs Kaserne, TMO 4th Trans.	1	1				
Rhein/Main, AATCO	1	1				
Rhein/Main, ACA, USAFE	1	1				
Rhein/Main, MASS MAC	1					1
Total	12	8	1	2		1

Location	Tot	HCp	Dis	Por	Rel	Bl
Ramstein						
Stage 1 Sites						
Ramstein, HQ USAFE	1	1				
Ramstein, USAFE TRAC	1	1				
Ramstein, 322nd ALD, MAC	1	1				
Ramstein, MASS MAC	1					
Ramstein, ACA, USAFE	1	1				
Ramstein, HQ USAFE OSC	1	1				
Ramstein, AATCO, 4th Trans.	1	1				
Kaiserslautern, HQ 2nd Mov.	1	1				
Kaiserslautern, TMO 4th Trans.	1	1				
Kaiserslautern, 37th Trans. Group	1	1				
Kaiserslautern, 21st SPT COM	1			1		
Kauserslautern, MOTCA, USAREUR	1	1				
Idar Oberstein, TMO 4th Trans.	1	1				
Zweibruecken, 200th TAMMC	1	1				
Total	14	12		1		
Heidelberg (HQ USAREUR)						
Stage 1 Sites						
Heidelberg, HQ USAREUR	3	2		1		
Mannheim, TMO	1	1				
Mannheim, MTMC TTGE	1	1				
Nuremberg TMO	2	1		1		
Total	7	5		2		

Location	Tot	HCp	Dis	Por	Rel	Blk
Stuttgart (Patch Barracks)						
Stage 1 Sites						
HQ USEUCOM	3	1	1	1		
TMO Stuttgart	1	1				
Total	4	2	1	1		
Naples (NAV SUP ACT)						
Stage 2 Sites						
ASCOMED	2	1	1			
COMSERVFOR 6th FLT	1	1				
COMFAIRMED	1	1				
COMSUBGRU 8	1	1				
Air Terminal NAV SUP ACT	2	1	1			
Water Terminal NAV SUP ACT	1	1				
ATOC	1	1				
Port Water OPNS	1	1				
NAV SUP ACT Supply	1	1				
NAV SUP ACT ADP	1	1				
MSCO	1	1				
Gaeta, COM 6th FLT	1	1				
Stage 2 Subtotal	14	12	2			
Stage 3 Sites						
Coltano, Camp Darby MTMC Terminal	1				1	
Sardinia, La Maddalena	1				1	
Stage 3 Subtotal	2				2	
Total	16	12	2		2	

Location	Tot	HCp	Dis	Por	Rel	Blk
Sigonella (NAF-SIG)						
Stage 2 Sites						
Air Terminal	2	1	1			
NAF ATOC	1	1				
NAF Supply	1	1				
NAF ADP	1	1				
VR-24	1	1				
Total	6	5	1			
Rota (NAV-STA Rota)						
Stage 2 Sites						
Naval Air Terminal	3	2	1			
ATOC	1	1				
VR-24	1	1				
Port (water) OPNS	1	1				
NAV STA Supply	1	1				
NAV STA ADP	1	1				
Cadiz, MTMC TTU	1	1				
Stage 2 Subtotal	9	8	1			
Stage 3 Sites						
Lisbon, MTMC Outport	1				1	
Torrejon, ACA USAFE	1				1	
Stage 3 Subtotal	2				2	
Total	11	8	1		2	

Location	Tot	HCp	Dis	Por	Rel	Blk
Athens (MTMC TTU, Piraeus)						
Stage 3 Sites						
Piraeus, MTMC TTU	1	1				
Crete, Souda Bay	1				1	
Total	2	1			1	
Istanbul (MTMC TUSLOG)						
Stage 3 Sites						
Istanbul, MTMC TUSLOG	1	1				
Izmir, MTMC TUSLOG	1				1	
Iskenderun, MTMC TUSLOG	1				1	
Incirlik, ACA USAFE	1				1	
Cakmakli, 67th DET.	1				1	
Total	5	1			4	

Grand Total of Terminals	Tot	HCp	Dis	Por	Rel	Blk
Stage 1	69	44	10	10		5
Stage 2	29	25	4			
Stage 3	16	5	2		9	
Total	114	74	16	10	9	5

7. APPENDIX B -- Traffic Survey Data

\RECD SEND\	LON	RDM	BRM	FRK	RAM	HDL	STT	NAP
LON	90 0 90				20 0 25			75 0 100
RDM			9 0 0	18 0 0	18 0 36	15 0 0	3 0 0	
BRM(*)		46 50 50	341 153 107	* 91 40 20	* 51 0 45	* 44 5 0	* 6 0 0	
FRK	18 0 42	27 123 99	39 213 300	24 432 1917	111 276 511	36 93 450	3 36 132	
RAM	20 0 20	0 0 35	15 21 61	157 12 235	60 282 712	775 18 67	15 0 30	60 0 60
HDL		0 21 0	0 96 66	12 159 189	19 156 178	0 84 45	0 3 48	
STT		0 21 0	0 96 24	0 156 90	15 114 119	0 90 150	0 0 0	
NAP	75 0 110				60 0 60			
SIG	100 0 65							2000 60 400
ROT	75 0 80				10 0 20			900 60 400
IST					20 0 30			
USA	40 0 75				125 0 25			90 0 120

* Add 500 msg/mo to BRM-FRK and to BRM-RAM; add 80 msg/mo to BRM-HDL and to BRM-STT...distribution of AUTODIN/Telex/Voice is unavailable.

SIG	ROT	IST	ATH	USA	TOT	/F /SEND
15	20			15	235	
0	0			0	0	LON
20	25			35	295	
				12	75	
				0	0	RDM
				0	36	
				15	594	
				0	248	BRM
				10	232	
				27	285	
				0	1173	FRK
				261	3712	
	10	20		825	1957	
	0	0		0	333	RAM
	20	30		25	1295	
					31	
					519	HDL
					526	
					15	
					477	STT
					383	
1500	1500			100	3235	
60	60			0	120	NAP
400	400			30	1000	
	90			60	2250	
	60			0	120	SIG
	600			30	1095	
150				120	1255	
60				0	120	ROT
400				300	1200	
					20	
					0	IST
					30	
90	120				465	
0	0				0	USA
100	100				420	

				AUTODIN	10417	
				TELEX	3110	
				VOICE	10224	

8. APPENDIX C -- Peak Traffic Estimation Methodology

We can estimate the number of bits per second that a node sends to its associated EM host, the number of bits per second that a node receives from its EM host, and the intra-EM host traffic in bits per second, during the busy hour. The estimates are based on the number of messages per month composed by users at the network nodes. We assume that there is one busy hour for the entire MINET (conservative overlapping of all users' busy periods) during which one quarter of the daily traffic is sent. We begin by estimating the number of subnetwork packets, of various types, that are transmitted to and from an EM host in order to compose and to read an EM message. These estimates are based on assumptions regarding typical user interactions required to compose and read EM messages. Since each user terminal is connected to a local TAC, the line-at-a-time behavior of the TAC is taken into account when calculating the number and the length of the subnetwork packets. A consideration of the data-carrying packets, as well as subnet acknowledgement packets (called "RFNMs" in the ARPANET) and TCP acknowldegments produces the following traffic estimates, which are in units of subnetwork packets:

Number of packets generated by the EM operation	short pkts	medium pkts	long pkts	TCP acks	subnet acks
msg compose, to host:	34	30	0	36	100
msg compose, from host:	6	30	0	64	100
msg read, to host:	3	0	0	5	8
msg read, from host:	2	0	17	3	8

Since we are assuming that the number of messages read is three times the number of messages composed, the total network traffic due to EM messages between a user and his EM host, for each of these five packet types, will be:

(# msgs composed) X (pkts/msg composed + 3 X (pkts/msg read)).

We can refer to the composition of one message and the reading of three messages as a "canonical transaction", so the total network traffic, in terms of each of these five packet types, can be expressed as:

(# msgs composed) X (pkts/canonical transaction).

By adding the "composed" and three times the "read" components of to-host traffic, and adding the "composed" and three times the "read" components of from-host traffic, we get:

Total packets per canonical transaction	short pkts	medium pkts	long pkts	TCP acks	subnet acks
pkts to host:	43	30	0	51	124
pkts from host:	12	30	51	73	124

We can convert from units of packets per canonical transaction to units of packets per second during the busy hour as shown:

$$\begin{array}{l} \text{pkts per} \\ \text{second in} \\ \text{busy hour} \end{array} = \frac{\text{packets}}{\text{canon trans}} \times \frac{\text{canon trans}}{\text{mo}} \times \frac{\text{1 mo}}{\text{20 day}} \times \frac{\text{1 day}}{\text{4 bsy hr}} \times \frac{\text{1 bsy hr}}{\text{3600 sec}}$$

By assumption, the number of canonical transactions per month is the number of messages composed per month. Thus,

$$\begin{array}{l} \text{pkts per} \\ \text{second in} \\ \text{busy hour} \end{array} = \frac{\text{msg comp}}{\text{mo}} \times \frac{\text{packets}}{\text{canon trans}} \times \frac{\text{1 mo}}{\text{288000 sec}}$$

We can convert packets to bits by estimating packet lengths. The short, medium, long, TCP ack, and subnet ack packets are assumed to be 544, 1000, 1198, 520, and 232 bits in length, respectively. The short, medium, and long packet lengths include the subnetwork and TCP header overhead. By augmenting the equation above to include the bits per packet estimates, we can express the traffic in units of bits per second during the busy hour. Multiplying these factors together and then summing across the five packet types, in both the to-host and from-host directions, yields:

$$\begin{array}{l} \text{Bits/sec} \\ \text{to host} \\ \text{in busy hr} \end{array} = 0.375 \times \frac{\text{\# msg comp}}{\text{mo}} \qquad \begin{array}{l} \text{Bits/sec} \\ \text{from host} \\ \text{in busy hr} \end{array} = 0.567 \times \frac{\text{\# msg comp}}{\text{mo}}$$

Intra-EM host traffic is modelled as a message read; EM host "A" sending a message to EM host "B" is treated as "B" reading a message from "A". Starting with the the number of packets of each type generated by the "message read" operation, reversing the "from" and "to" directions, and multiplying by the above conversion factors and then summing across the five packet types, we obtain:

$$\begin{array}{l} \text{Bits/sec} \\ \text{to host} \\ \text{in busy hr} \end{array} = .089 \times \frac{\text{\# msg comp}}{\text{mo}} \qquad \begin{array}{l} \text{Bits/sec} \\ \text{from host} \\ \text{in busy hr} \end{array} = .023 \times \frac{\text{\# msg comp}}{\text{mo}}$$

The only input data needed to produce these estimates are the number of messages composed per month at each network node.

9. APPENDIX D -- Response Time Analysis Methodology

We begin by assuming that the major components of propagation delay in the MINET will come from transmission, propagation and queueing delays associated with the 9.6 Kbps trunk lines that connect IMPs. The same IMPs are used in the ARPANET with 50 Kbps trunks connecting them, so we expect that they will be able to handle MINET traffic without becoming CPU-bound. We will include a small fixed component of delay for the time associated with processing in IMPs and TACs. We model each line in the MINET as a server with fixed service rate of 9.6 Kbps. The traffic between sites reported to us is distributed over each line using a min-hop fixed path routing algorithm. The routing algorithm used in the ARPANET can be simulated under steady state conditions using a min-hop strategy. The strategy must maintain single deterministic paths from all points to all other points with no conflicts in simulated routing tables. None of the analysis is valid for periods of traffic or network transition. The min-hop routing produces a logical network map with steady state trunk utilizations as shown in Figure D-1. This allows us to estimate the utilization of each user path in the MINET. Using standard queueing theory formulas for M/G/1 systems, we can then estimate the time packets spend waiting to be transmitted over each of the 9.6 Kbps lines.

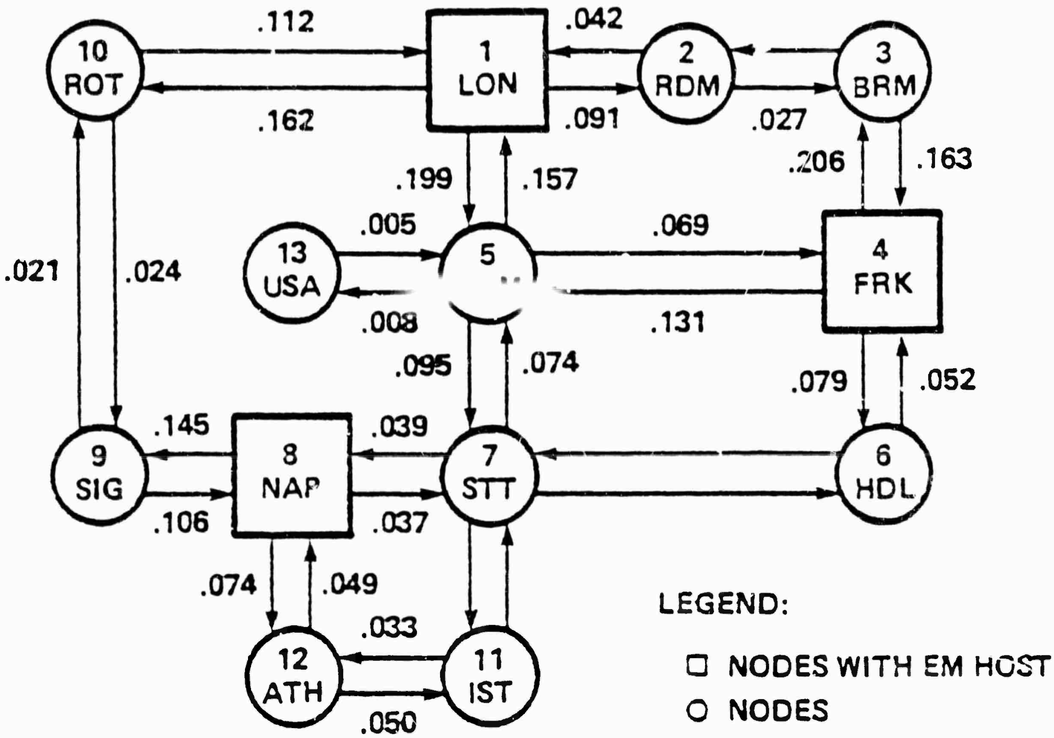

Figure D-1 MINET Logical Map Under Normal Conditions

The values that were used in our calculations are shown below. All numbers are one-way delays expressed in milliseconds. Where values differ significantly for heavily loaded and lightly loaded network conditions (normal traffic and four-times normal traffic respectively), the value shown in parentheses applies to the heavily loaded condition.

TAC processing and queueing delays	10
End IMP (to which host or TAC is attached)	50 (75)
Tandem IMP (store and forward only)	10 (20)
Trans time for packets TO host (short)	57
Trans time for packets FROM host (medium)	104

Other factors used include propagation delay of 20 us/km where the distance in km is the straight line distance between two IMPs (this is 1/3 the speed of light times a factor of 2 to take into account actual circuit route length). The M/G/1 queueing model predicts queueing delays of xp(1+C)/(2*(1-p)) where x is the service time of the lines (based on 9.6 Kbps processing of the various packet lengths), p is the utilization of the line, and C is the square of the coefficient of packet variation (computed to be approximately .45 for traffic going towards a host and .12 for traffic going towards a terminal). By applying these assumptions to each node under varying traffic and failure conditions, we have computed estimated typical transmission delay times.